

CHAIN-LEVEL STRING TOPOLOGY OPERATIONS

GABRIEL C. DRUMMOND-COLE, KATE POIRIER, AND NATHANIEL ROUNDS

ABSTRACT. We construct a space of string diagrams, which are a type of fatgraph with some additional data, and show that there are string topology operations on the chains of the free loop space of a closed Riemannian manifold which are parameterized by the chains on the space of string diagrams. These operations are shown to recover known structure on homology of the free loop space.

CONTENTS

Introduction	1
1. Conventions and definitions	7
2. String diagrams	12
3. Straightening a string diagram	23
4. The heart of the string topology construction	26
5. The push-pull map and the string topology construction	31
6. Patching cohomology classes over the space of oriented string diagrams	33
7. The diffuse intersection class	43
8. Connections to previous work	48
Appendix A. Straightening of short-branched trees	58
References	66

INTRODUCTION

String topology on a manifold M is the study of natural operations on the free loop space LM , the space of continuous maps from circle to M . These operations on LM are defined using the algebraic and geometric topology of M itself.

This work was supported by IBS-R003-D1.

This work was partially supported by NSF RTG grant DMS-0838703.

Historically, string topology operations were constructed as maps between various tensor powers of the homology groups of LM . Usually these operations involve a shift in degree, which we suppress throughout this introduction. These homology-level operations can be parameterized by spaces of graphs; the spaces of graphs parameterizing operations discovered so far are closely related to moduli spaces of Riemann surfaces.

In this paper, we define string topology operations that are maps between the tensor powers of the singular *chains* of LM , rather than the homology groups of LM . These chain-level operations induce operations at the level of homology, including many of the previously constructed operations.

This paper constitutes one step in a program to extend the construction of string topology operations to the most general possible setting, along the way employing those spaces of operations that yield the finest possible invariants.

Background. This history of string topology is necessarily incomplete and is geared toward the contents of this paper.

Chas and Sullivan introduced string topology operations by defining a loop product on the homology groups of the loop space

$$H_*(LM) \otimes H_*(LM) \rightarrow H_*(LM)$$

which gives $H_*(LM)$ the structure of a commutative associative algebra [CS99].

A chain in LM determines a chain in M by evaluation of loops at their basepoints. For a pair of chains in LM whose chains of basepoints intersect transversally in M , Chas and Sullivan defined their chain-level loop product by concatenating loops along the intersection locus of basepoints in M . Since it is not true that any two chains in LM determine transversal chains in M , this chain-level product is only partially defined on $C_*(LM) \otimes C_*(LM)$. However, for each pair of homology classes in $H_*(LM)$, Chas and Sullivan choose pairs of representative cycles whose chains of basepoints in M do intersect transversally in order to give their loop product, which is fully defined on $H_*(LM) \otimes H_*(LM)$.

Later, Cohen and Jones took a different approach, using homotopy theory to generate an alternate definition of Chas and Sullivan's loop product [CJ02]. This allowed them to avoid explicitly working with transversality, which is notoriously delicate. The space $LM \times_M LM$ of pairs of loops in M whose basepoints coincide is a finite codimension subspace of the infinite dimensional manifold $LM \times LM$ of pairs of loops in M . Because of this Cohen and Jones were able to use a Pontryagin–Thom construction to define a wrong-way map from $H_*(LM \times LM)$ to $H_*(LM \times_M LM)$, which is a composition of two maps. The first map in the composition is induced on homology by the map of spaces which takes $LM \times LM$ to the Thom space of a tubular neighborhood of $LM \times_M LM$ in $LM \times LM$. The second is the Thom isomorphism from the homology of this Thom space to the homology of $LM \times_M LM$. Cohen and Jones realized the loop product by composing this wrong-way map with the map induced on homology by concatenating loops whose basepoints coincide.

Chas and Sullivan also showed that the loop product is, in fact, part of a richer structure on $H_*(LM)$, namely that of a Batalin–Vilkovisky algebra. The Batalin–Vilkovisky operator Δ , which is a square zero unary second-order differential operator with respect to the loop product, is induced by the S^1 action on LM given by rotating loops. Additionally, Chas and Sullivan used the loop product to define the string bracket, which gives $H_*^{S^1}(LM)$, the S^1 -equivariant homology of LM , the structure of a Lie algebra. They later used similar methods to define the string cobracket, which, together with the string bracket, gives $H_*^{S^1}(LM, M)$, the S^1 equivariant homology of LM relative to the subspace of constant loops, the structure of an involutive Lie bialgebra [CS04].

Chas and Sullivan’s operations inspired Cohen and Godin to develop a coherent set of more general operations [CG04] following the method of Cohen and Jones. Cohen and Godin defined a kind of fatgraph called a Sullivan chord diagram and added additional structure to these diagrams to define a *space* of marked metric chord diagrams.

For each point in their space, they defined an operation of the form

$$H_*(LM)^{\otimes k} \rightarrow H_*(LM)^{\otimes \ell}$$

for some k and ℓ determined by the diagram. They also showed that this operation depends only on the path component of the point in their space. In other words, they described natural operations on the homology of the loop space parameterized by the zeroth homology of the space of marked metric chord diagrams. Chataur later showed that there are natural operations on the homology of the loop space parameterized by the *higher* homology of the space of marked metric chord diagrams [Cha05].

There is an embedding of Cohen and Godin’s space of marked metric chord diagrams into the moduli space \mathcal{M} of Riemann surfaces with parameterized boundary. Godin later extended Cohen–Godin’s construction to provide natural operations on the homology of the loop space parameterized by the homology of \mathcal{M} [God07a]. Kupers recently used a different model of \mathcal{M} originally defined by Bødigheimer to recover Godin’s structure [Kup13, Bød06]. In fact, Cohen–Godin, Chataur, Godin, and Kupers each assembled their collections of operations into an algebraic structure, in each case some kind of field theory.

Tamanoi [Tam10] showed that many of Cohen and Godin’s operations are trivial. He also showed that Godin’s operations coming from homology classes of \mathcal{M} which are in the image of the stabilization map are trivial [Tam09]. The full implications of these facts for Chataur and Godin–Kupers’ higher homology operations are still unclear, but this presents an interesting question because not much is known about the unstable homology of \mathcal{M} . Given a homology class of \mathcal{M} , is there a manifold M for which the corresponding string topology operation is nontrivial? If the answer is yes, then the homology class is an unstable class. As far as the authors know, so far no unstable homology classes have been found by these means.

Other lines of research have focused on algebraic models of the free loop space rooted in Hochschild homology [TZ06, Mal11]. Some of this research has also used spaces of chord diagrams. This Hochschild approach is not taken in this paper and we will not discuss this perspective in detail.

Compactified and chain-level operations. A general precept of homotopical algebra is that natural operations on the homology of something usually reflect a richer structure at the chain level, and that the passage to homology usually loses information. Following this precept, a natural conjecture is that all the algebraic structure on the homology $H_*(LM)$ of the free loop space discussed above is the shadow of an algebraic structure on the chains $C_*(LM)$. A second conjecture is that this algebraic structure on the chains is strictly richer than that on homology. A realization of the first conjecture is the subject of this paper; at this point the status of the second conjecture is still open, as far as we know.

Another natural goal is to describe the universal space of string topology operations. With this goal in mind, we expect that there are natural operations on the chains of the free loop space parameterized by the chains on a *compactification* of M . We expect that these operations recover those parameterized by the non-compact spaces used by [CG04, Cha05, God07a, Kup13] and also incorporate operations and relations not evident in this previous work.

These considerations were discussed by Sullivan in his survey [Sul05]. There he outlined natural operations on a chain complex computing the S^1 -equivariant homology of the loop space parameterized by a compactification of a related moduli space of Riemann surfaces. Following Sullivan, there have been multiple descriptions of chain-level operations and compactified spaces of operations. The second author, in her thesis [Poi10], described a compactified space of operations and a more detailed description of the operations on this equivariant chain model. The second and third author wrote a preprint [PR11] describing a non-equivariant version; this paper is the natural outgrowth of that preprint but uses new tools to give what we believe is a more natural and comprehensive presentation. Recently, Irie defined a model for chains on the free loop space to deal with transversality issues [Iri14, Iri15]. Using this model, he was able to define chain-level operations parameterized by a space of graphs called decorated cacti. However, this model does not seem to be well-suited to operations with more than one output or corresponding to surfaces of higher genus. Hingston and Wahl [HW14] have also announced results similar to ours.

Contents of this paper. Following Sullivan [Sul05], our chain-level construction deals with transversality issues by using short geodesic segments, and generalizations of such segments, to join points which are nearby in M , but which may not actually coincide. It combines aspects of both the original Chas–Sullivan construction for transversally intersecting chains and the Cohen–Jones Pontryagin–Thom construction. This geodesic segment technique is our replacement for the concatenation of loops that literally intersect; our use of this technique is a fundamental difference between our methods and those of Chas–Sullivan and Cohen–Jones.

In this paper we define a version of chord diagrams called *string diagrams*, which generalizes the definitions of string diagrams appearing in [PR11] and [Poi10]. Roughly, string diagrams are metric graphs obtained by attaching leaves of trees to circles.

We denote the space of string diagrams by \mathcal{SD} . Each path component $\mathcal{SD}(\chi, k, \ell)$ of \mathcal{SD} has the structure of a finite cell complex so is compact.

We conjecture that the space of string diagrams is homotopy equivalent to the moduli space of Riemann surfaces with parameterized boundary. There is a natural equivalence relation on the space of string diagrams and we further conjecture that our space modulo equivalence has the homotopy type of a known compactification. See the next subsection for details.

The main theorem of the paper is the following.

Main Theorem. *There are chain-level string topology operations*

$$C_*(\mathcal{SD}(\chi, k, \ell)) \otimes C_*(LM)^{\otimes k} \rightarrow C_*(LM)^{\otimes \ell}.$$

which realize diffuse intersection string topology and induce the classical string topology operations of [CS99] and [CG04].

The operations are defined in Section 4 and they are shown to recover previous work in Section 8. Dealing with transversality issues requires care and results in some technical definitions.

The most important technical accomplishments in the paper are the following.

- (1) We construct an appropriate *diffuse intersection class* representative that we can cap chains with to model intersection. This involves carefully patching together Thom classes and ensuring that the output class is well-defined and unique. This diffuse intersection class plays the role of the Thom class in our version of the Pontryagin–Thom construction.
- (2) We define a generalized geodesic construction \heartsuit , so-called because it is the heart of our string topology construction. The map \heartsuit takes a string diagram and a collection of loops in M satisfying a closeness condition dependent on the string diagram and outputs a map from the string diagram to M . This must be done coherently in families as the string diagram varies in \mathcal{SD} .

Carefully defining these two constructions takes up much of this paper.

We show that our operations induce Cohen–Godin’s operations and Chas–Sullivan’s BV algebra on homology. We expect that our operations should also induce the structures of Chataur, Godin, and Kupers as well.

This paper is organized as follows.

In Section 1 we introduce our conventions and preliminary definitions which will be used for the definition of string diagrams given in Section 2. Also in Section 2, we discuss the space of string diagrams and show that it has the structure of a cell complex.

In Section 3 we show how to map a particular subgraph of a string diagram into a simplex by first mapping its leaves to the vertices. This *straightening* map is defined in terms of a *straightening* map for trees in Appendix A which maps a tree into a simplex by first mapping its leaves to the vertices. The straightening map for string diagrams is the first map in a composition of two maps defining the map \heartsuit . Section 4 is devoted to the definition of the map \heartsuit , our generalized geodesic construction.

In Section 5 we fix an arbitrary cochain W on the domain S of the map \heartsuit and use it to define a *push-pull* map $\mathcal{ST}_W : C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k} \rightarrow C_*(LM)^{\otimes \ell}$. When the cochain W is a cocycle representing the *diffuse intersection class*, we call \mathcal{ST}_W a *string topology construction*.

The diffuse intersection class is the pullback under an *evaluation map* of a *global Thom class over \mathcal{SD}* . In Section 6 we prove Theorem 6.18 to describe how to assemble cohomology classes which are indexed by cells of \mathcal{SD} to get *global classes over \mathcal{SD}* . In Section 7 we assemble the global Thom class over \mathcal{SD} and define the evaluation map in order to define the diffuse intersection class.

In Section 8 we show that our string topology operations recover Cohen–Godin’s operations and Chas and Sullivan’s BV algebra structure on $H_*(LM)$.

Future directions. There are several points we have left unaddressed in this paper.

First of all, there is the question of composition. Cohen and Godin described a sort of composition map on marked metric chord diagrams. It would perhaps be better to call it a composition relation, as it is not defined for all pairs and is not always unique when it is defined. They showed that this composition map induces a well-defined associative map on zeroth homology groups and that their string topology construction respected composition.

In our context too there is a cognate composition map. Our composition map is fully defined and uni-valued. However, it is not strictly associative but only associative up to homotopy. Similarly, the chain-level string topology construction does not respect composition on the nose, but rather only up to homotopy. We could rectify this algebraic structure or use colored or ∞ -properads to deal with this problem, but as this paper is already long and technical we decided instead to defer these questions.

Another question is that of the homotopy type of the space of string diagrams. As discussed previously, we have the following conjecture:

Conjecture A. The space of unoriented string diagrams is homotopy equivalent to the moduli space of Riemann surfaces with parameterized boundary.

There is an equivalence relation on the space of string diagrams, not discussed in this paper. Our constructions pass to the quotient by the equivalence relation. We also formulate the following conjecture:

Conjecture B. The space of unoriented string diagrams modulo equivalence is homotopy equivalent to Bødigheimer’s harmonic compactification of the moduli space of Riemann surfaces with parameterized boundary [Böd06].

This compactification has appeared before in string topology, in the work of the second author as well as Kupers [Poi10, Kup13].

There are a few other avenues of research that should at this point be straightforward, if technical, generalizations of the techniques used in the current paper.

First, Godin's field theory in fact gives operations for both loops and paths in M ; it is an *open-closed* theory. While in this paper we focus purely on closed loops, we expect the methods that produce our string topology operations to generalize to an open-closed chain-level theory.

Second, we expect that our operations induce chain-level operations which recover Chas and Sullivan's involutive Lie bialgebra on reduced equivariant homology, which would enable us to realize Sullivan's outline from [Sul05] in more detail.

Finally, we expect that we should be able to use cellular chains on our space of string diagrams to define operations on the Hochschild-homological models of the free loop space.

Acknowledgements. The authors gratefully acknowledge many helpful discussions with Dennis Sullivan. The second author would like to thank Thomas Tradler for helpful conversations and thank the IBS Center for Geometry and Physics in Korea for a lovely and conducive working environment.

1. CONVENTIONS AND DEFINITIONS

1.1. Conventions. A *graph* is a quadruple (V, H, s, ι) , where V is a finite set of *vertices*, H is a finite set of *half-edges*, s is a surjective map from H to V (we call $s(h)$ the *source* of h), and ι is a fixed-point free involution of H . We call an orbit of ι an *edge*. The *half-edge set* of a vertex v is the set of half-edges with source v . The *valence* of a vertex is the cardinality of its half-edge set. A univalent vertex is called a *leaf*; a vertex which is not a leaf is called an *internal vertex*. A *leaf half-edge* (respectively an *internal half-edge*) is a half-edge whose source is a leaf (respectively an internal vertex). An *external edge* is an edge containing a leaf half-edge; an *internal edge* is an edge not containing a leaf half-edge.

A graph is *connected* if it admits no partition into two disjoint nonempty subgraphs. A *component* of a graph is a maximal connected subgraph. A *cycle graph* is a nonempty connected bivalent graph. A *tree* is a nonempty connected graph with no cycle subgraph. A *segment* is a tree with one edge. The *Euler characteristic* of a graph is the difference between its number of vertices and its number of edges.

A *cyclic order* on a finite set is a permutation of that set which is a single cycle. We shall call a set equipped with a cyclic order a *cycle*. A *fatgraph* Γ is a graph G together with a cyclic order of the half-edge set of each vertex of G .

A *pseudometric (fat)graph* is a (fat)graph together with a non-negative length function ℓ on the set of edges. Let Γ be a pseudometric (fat)graph. The *length* of Γ is the sum of the lengths of all the edges of Γ .

Let e be an edge of a graph $G = (V, H, s, \iota)$ which is not a segment component. *Contracting the edge e* produces a graph G/e which has half-edge set $H - e$ and vertex set the quotient of V where the sources of e are identified. The source map

and involution for G/e are induced by those for G . If G has a fat or pseudometric structure then G/e inherits a fat or pseudometric structure from G ¹.

An *oriented edge* \vec{e}_h of a graph consists of a choice of order $(h, \iota(h))$ for the two half-edges that make up the edge $e_h = \{h, \iota(h)\}$. The source of the first half-edge is called the *source* of the oriented edge and the source of the second half-edge is called the *target* of the oriented edge. The set \tilde{E} of oriented edges of a graph is in bijective correspondence with its set of half-edges. Hence, \tilde{E} inherits the involution ι from the set of half-edges and $\iota(\vec{e}_h)$ is equal to $\vec{e}_{\iota(h)}$.

The set \tilde{E} also forms a double cover of the set E of edges of a graph, where the map forgets the ordering of the constituent half-edges. If G is a pseudometric graph, then we define the length of an oriented edge to be equal to the length of its underlying unoriented edge.

A *boundary cycle* of a fatgraph is a cycle of half-edges h_1, \dots, h_n so that h_{i+1} follows $\iota(h_i)$ in the cyclic order of the half-edge set of $s(h_{i+1})$. See Figure 2.

Remark. There is a standard construction which produces an oriented surface with boundary from a fatgraph using the fatgraph structure to thicken it [Str84]. The boundary cycles of a fatgraph correspond precisely to the boundary components of this surface. In what follows we do not use the thickening explicitly, though it is sometimes included in the figures.

A *marking* of a boundary cycle C of a fatgraph Γ is a leaf of Γ whose leaf half-edge appears in C . If G is pseudometric, we require that the length of the external edge is zero. A *partially marked fatgraph* is a fatgraph such that each boundary cycle contains at most one marking. A *marked fatgraph* is a fatgraph such that each boundary cycle contains exactly one marking.

1.2. Preliminary definitions. Next, we collect several less standard definitions for graphs that we will use in our main definition of Section 2.

Informally, a *vertex explosion* of a graph pulls apart a graph at a specified vertex, making an n -valent vertex into n distinct leaves.

Definition 1.1. Let $G = (V, H, s, \iota)$ be a graph and let v a vertex of G . The *vertex explosion* of G at v is a new graph whose vertices are

$$V - \{v\} \sqcup (\{v\} \times s^{-1}(v)).$$

The half-edge set and involution are not changed; the source map takes h in $s^{-1}(v)$ to $v \times h$ and is otherwise unchanged.

Remark. Up to canonical isomorphism, vertex explosion at different vertices commutes so we can unambiguously take a vertex explosion at a set of vertices.

We can also do a kind of partial vertex explosion.

¹Many definitions do not allow contraction of edges whose sources coincide or edges of positive length. In practice we will never contract such edges.

Definition 1.2. Let $G = (V, H, s, \iota)$ be a graph. Let v a vertex of G and let S be a proper subset of $s^{-1}(v)$. The *pruning* of G at v along S is a new graph whose vertices are

$$V \sqcup (\{v\} \times S).$$

The half-edge set and involution are not changed; the source map takes h in S to $v \times h$ and is otherwise unchanged.

Figure 1 shows an example of pruning a tree at an internal vertex along a single half-edge.




FIGURE 1. A tree T and the pruning of T at v along the half-edge h .

Remark. There is a canonical quotient map from any vertex explosion or pruning of G to G itself. If G is fat and/or pseudometric, then the same structure is induced on a vertex explosion or pruning of G .

Let U -graphs stand for either graphs, fatgraphs, pseudometric graphs, or pseudometric fatgraphs. We will tend to abuse notation by specifying only the data of the underlying graph of a U -graph, eliding the fat or pseudometric structure until it is needed.

There is a *pseudometric realization* $|\cdot|$ functor from pseudometric graphs to metric spaces as follows. Consider the metric space

$$V \sqcup \bigsqcup_{\vec{e} \in \tilde{E}} ([0, \ell(\vec{e})] \times \{\vec{e}\})$$

(where the distance between two points in different components is defined to be the length of G). The space $|G|$ is defined as the quotient metric space by the following relations. For the oriented edge \vec{e} , $(0, \vec{e})$ is identified with the source of \vec{e} and $(\ell(\vec{e}), \vec{e})$ is identified with the target of \vec{e} . Further, (t, \vec{e}) is identified with $((\ell(\vec{e}) - t), \iota(\vec{e}))$. See Figure 2.

Note that in the case that the pseudometric graph has no length zero edges, the pseudometric realization is homeomorphic to the usual geometric realization of a graph. In general, the pseudometric realization is homeomorphic to a quotient of the usual geometric realization obtained by contracting all length zero edges. This




FIGURE 2. A pseudometric fatgraph G , its pseudometric realization $|G|$, a marked boundary cycle C , and the map ∂_C from a standard circle into $|G|$.

means that if e is an edge of length zero of the pseudometric graph G , there is a canonical identification of $|G|$ with $|G/e|$.

We will have need for the following graph model of the circle. A *fat lollipop* is a connected fatgraph which has an external edge between a leaf and a trivalent vertex and all other vertices bivalent. We call the boundary cycle containing the leaf half-edge the *marked boundary cycle*. See Figure 3.




FIGURE 3. A pseudometric fat lollipop of total length one G and its pseudometric realization $|G|$.

We would like to use a boundary cycle to define a map from the circle into the realization of a pseudometric fatgraph. To do this, we require exactly one marking of the boundary cycle. Our model for the *length L standard circle* is the quotient of the interval $[0, L]$ which identifies 0 with L . Linear scaling provides identifications of standard circles of different lengths. Let C be a boundary cycle of a pseudometric fatgraph G with exactly one marking v . There is a canonical map ∂_C from a standard circle into the pseudometric realization $|G|$ of G , defined as follows. Let C be the marked boundary cycle (h_1, \dots, h_n) where h_1 is the leaf half-edge of the leaf v and h_n is equal to $\iota(h_1)$. Let \vec{e}_i be the oriented edge $(h_i, \iota(h_i))$ with length ℓ_i . Let L be the sum of the lengths ℓ_i . If $L = 0$, ∂_C is the constant map to the image of the leaf v in the pseudometric realization $|G|$ of G . If $L > 0$, ∂_C is induced

by a map from $[0, L]$ to $|G|$ which maps the subinterval $[\ell_1 + \dots \ell_{i-1}, \ell_1 + \dots \ell_i]$ to the image of $[0, \ell_i] \times \vec{e}_i$ in $|G|$ by translation on the first factor. This map sends L to the image of v in $|G|$, so it determines the map ∂_C from the length L standard circle to $|G|$. See Figure 2. We will also use the map $\bar{\partial}_C$ from the length L standard circle to $|G|$ which uses the same algorithm but follows the boundary cycle in the opposite order. Alternatively, this can be written $\bar{\partial}_C(t) = \partial_C(L - t)$.

We will be considering pseudometric fatgraph trees in service of our main definition, Definition 2.4. Later these trees will parameterize our diffuse intersections. For technical reasons it will be useful to have the distance between distinct leaves in such a tree be both bounded above and bounded away from zero (see Definition 1.6). We will achieve the bound above by imposing a length restriction on the trees we consider. We will achieve the bound away from zero by imposing a further length restriction on certain subtrees. This will require some setup.

Definition 1.3. The *leaf length* of a pseudometric fatgraph tree T is one less than the number of leaves of T .

The definition of leaf length may seem unmotivated, but see the remark following the next definition.

Definition 1.4. Let h be a half-edge of a tree T whose source is at least trivalent. The *branch* T_h is the unique maximal subtree of T containing h so that the vertex of h is a leaf in T_h . The *pollard* T^h is the unique maximal subtree of T which contains no edge of T_h .

Remark. The leaf length of a branch of T is the number of its leaves which are leaves in T .

Remark. We can canonically identify T_h with the component of the pruning of T at $s(h)$ along $\{h\}$ which contains h and T^h with the component of the same pruning which does not contain h . See Figure 4.




FIGURE 4. The branch T_h and pollard T^h of a tree T and half-edge h .

Recall that the *length* of a pseudometric graph is the sum of the lengths of its edges.

Definition 1.5. We call a branch of a pseudometric tree *prunable* if its length is equal to its leaf length.

Definition 1.6. A pseudometric tree T is *short branched* if:

- (1) The length of T is equal to its leaf length.
- (2) For every half-edge h of T whose source is at least trivalent, the length of the branch T_h is less than or equal to its leaf length.

We call a short-branched tree *non-degenerate* if it has no edges of length zero and no prunable branches. We call a short-branched tree *degenerate* otherwise.

Figure 5 shows examples of short-branched pseudometric trees and pseudometric trees that are not short branched.




FIGURE 5. Examples of short-branched trees and trees which are not short branched. Edge lengths are given. The subtrees made up of bold edges are examples of branches whose lengths are greater than or equal to their leaf lengths.

2. STRING DIAGRAMS

2.1. Definitions of string diagrams. The main two definitions of this paper are Definitions 2.1 (a combinatorial definition) and 2.4 (which adds continuous pseudometric data). In these definitions, we describe the type of graphs we will work with. Intuitively, we start with a collection of fat lollipops, $\{Q_i\}$. Then we glue on a collection of fatgraph trees $\{T_j\}$ along their leaves. The leaves of the first fatgraph trees are glued onto the fat lollipops, but successive trees can have some or all of their leaves glued onto trees that had been glued on previously as well. Finally, we glue on markings $\{L_i\}$ onto any boundary cycle which is not already marked. The rigorous presentation of Definition 2.1 does not follow this intuitive outline exactly, but captures the ideas behind it.

Definition 2.1. A *combinatorial string diagram* Γ is a connected marked fatgraph (also called Γ) with no bivalent vertices equipped with the following data:

- (1) collections of subgraphs $\{Q_1, \dots, Q_k\}$, $\{L_1, \dots, L_\ell\}$, and $\{\tilde{T}_j\}$ for j in some finite set so that each edge of Γ is contained in precisely one subgraph, and
- (2) a choice of a subset of internal vertices of each \tilde{T}_j as *fundamental*.

This data satisfies the following conditions:

- (1) Each Q_i is a fat lollipop whose inclusion into Γ preserves its marked boundary cycle,
- (2) each L_i is a segment precisely one of whose leaves is a leaf of Γ ,
- (3) no \tilde{T}_j contains a leaf of Γ ,
- (4) each internal vertex of Γ is either contained in precisely one Q_i or is disjoint from all Q_i and fundamental in precisely one \tilde{T}_j ,
- (5) the vertex explosion of each \tilde{T}_j at its non-fundamental vertices is a fatgraph tree, and
- (6) there is no cycle $\tilde{T}_{j_1}, \dots, \tilde{T}_{j_m} = \tilde{T}_{j_1}$ such that each \tilde{T}_{j_i} has a fundamental vertex which coincides with a non-fundamental vertex of $\tilde{T}_{j_{i+1}}$.

We denote the vertex explosion of \tilde{T}_j at its non-fundamental vertices by T_j .

See Figure 6 for a picture of a combinatorial string diagram.

Remark. We can recover \tilde{T}_j and its subset of fundamental vertices from the isomorphism class of T_j and the map

$$T_j \rightarrow \tilde{T}_j \rightarrow \Gamma.$$

That is, a vertex is fundamental in \tilde{T}_j if and only if it is internal in T_j . We will often specify the additional data necessary to turn a fatgraph into a combinatorial string diagram by specifying $T_j \rightarrow \Gamma$ rather than $\tilde{T}_j \subset \Gamma$. The final condition captures the idea from the intuitive outline that the T_j should be glued on according to a partial order on the set of trees $\{T_j\}$.

Notice that the leaves of the Q_i and the leaves of the L_i which are also leaves of Γ provide the markings of the boundary cycles of a combinatorial string diagram Γ .

Definition 2.2. We call the k boundary cycles C_1, \dots, C_k of a combinatorial string diagram Γ induced by the marked boundary cycles of Q_i *input boundary cycles*. Similarly, we call the other ℓ boundary cycles of Γ *output boundary cycles*.

Later, for the string topology construction, we will use combinatorial string diagrams to describe how a k -fold loop in a manifold may intersect—or nearly intersect—itsself. In order to describe the combinatorics of these intersection configurations we will use the components of the following fatgraph associated to a combinatorial string diagram.

Definition 2.3. Let Γ be a combinatorial string diagram. The *intersection graph* of Γ is the fatgraph $\hat{\Gamma}$ obtained as follows. Begin with the subgraph of Γ which is the union of all the \tilde{T}_j . Then perform vertex explosion on this subgraph at all vertices which are vertices of the Q_i subgraphs of Γ . See Figure 7.

Remark. The canonical map from T_j to Γ naturally factors through $\hat{\Gamma}$.

Definition 2.4. A *string diagram* is a combinatorial string diagram Γ equipped with a pseudometric structure such that

- (1) each Q_i has total length one, and
- (2) each T_j is short branched.

By abuse of notation, we also use Γ to denote the string diagram with underlying combinatorial string diagram Γ .

Definition 2.5. We call the disjoint union of k copies of the standard circle, equipped with the map given by the disjoint union of the canonical maps ∂_{C_i} to $|\Gamma|$, the *inputs* of Γ and denote the map by In_Γ . Similarly, we call the disjoint union of ℓ copies of the standard circle, equipped with the analogous map for output boundary cycles of Γ , the *outputs* of Γ .

Remark. The map from the inputs to $|\Gamma|$ is always a homeomorphism onto its image. The same need not be true for the outputs.

Definition 2.6. An *ordering* on a combinatorial string diagram is

- (1) an ordering on the set $\mathcal{T} = \{T_j\}$ and
- (2) an ordering on the disjoint union \mathcal{L} of the sets of leaves $L(T_j)$.

There is an evident action of the product of the symmetric groups of \mathcal{T} and \mathcal{L} on the set of orderings of a string diagram. There is a homomorphism from this group to $\mathbb{Z}/2\mathbb{Z}$ given by the product of the sign homomorphisms. An element in this product group is a pair of permutations; the element is in the kernel of the homomorphism to $\mathbb{Z}/2\mathbb{Z}$ if either both are even permutations or both are odd permutations.

Definition 2.7. An *orientation* on a combinatorial string diagram is a choice of ordering up to the action of the kernel of the homomorphism described above.

For a string diagram, an ordering or orientation means the respective structure on its underlying combinatorial string diagram.

Orientations will allow us to build Thom classes of products of diagonal maps of our manifold M in Section 7; they will not be used in a fundamental way until then.

2.2. The space of string diagrams. The rest of the section is devoted to describing the topology and CW structure on the set of (isomorphism classes of) oriented string diagrams.

Lemma 2.8. *Let Γ be a combinatorial string diagram. Let K be the set of string diagrams with underlying combinatorial string diagram Γ .*

Taking internal edge lengths defines an inclusion map from K to \mathbb{R}^E where E is the set of internal edges of Γ . The image of this inclusion is a bounded convex polytope (intersection of half-spaces) in \mathbb{R}^E .

Furthermore, passing to the boundary of this polytope corresponds to the following types of degenerations:

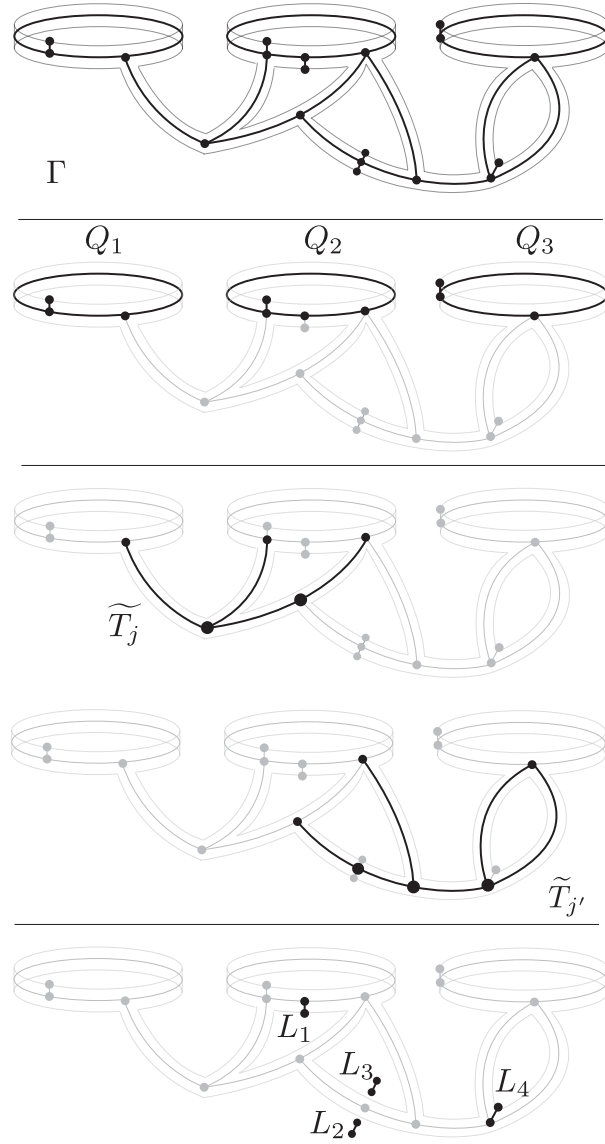


FIGURE 6. The data of a combinatorial string diagram Γ . Fundamental vertices are depicted as larger than non-fundamental vertices.

- (1) allowing edge lengths to shrink to zero, preserving the overall length of each Q_i and T_j and
- (2) allowing edge lengths to change so that T_j has a prunable branch.




FIGURE 7. A combinatorial string diagram Γ and its intersection graph $\hat{\Gamma}$.




FIGURE 8. Degenerations of the types (1)–edge lengths shrink to zero–and (2)–branches grow–listed in Lemma 2.8. Compare to Figure 11, where degenerate string diagrams at the boundary of the cell are identified with nondegenerate string diagrams in a different cell.

Proof. The set K has parameters corresponding to the lengths of edges in Γ . Since external edges always have length zero, the set of possible edge-length parameters form a subspace of $\mathbb{R}_{\geq 0}^E$. The lengths are subject to the conditions:

- (1) each Q_i has total length one,
- (2) each T_j has total length equal to its leaf length, and
- (3) each branch of a T_j has total length at most its leaf length.

Each condition in (1) and (2) corresponds to a hyperplane in \mathbb{R}^E . The intersections of all these hyperplanes with $\mathbb{R}_{\geq 0}^E$ is a bounded convex polytope, in fact a product of simplices. Each condition (3) corresponds to a half-space in \mathbb{R}^E . The intersection

of these half-spaces with the convex polytope is still a convex polytope. The faces of this polytope correspond to edges of length zero and prunable branches.

□

By abuse of notation, we use K to denote the set of string diagrams with underlying combinatorial string diagram Γ , given the subspace topology and polytope structure from the inclusion of the lemma.

The space of oriented string diagrams with a fixed underlying combinatorial diagram is the disjoint union of two copies of such a polytope. We will define a cell complex built by assembling the oriented version of these polytopes below. In preparation, we discuss two operations on string diagrams that we will use to define the attaching maps.

Definition 2.9. Let Γ be a string diagram with e an internal edge of length zero. We define a string diagram whose underlying pseudometric fatgraph is Γ/e , the fatgraph Γ with the edge e contracted as follows. We will specify subgraphs Q_i and L_i along with pseudometric fatgraph trees T_j equipped with a map to Γ/e as in the remark following Definition 2.4.

- (1) If e is an edge of Q_i , then contract e in Q_i and let all data be induced.
- (2) Similarly, if e is an internal edge of T_j or an external edge of T_j whose internal vertex is bivalent in T_j , then contract e in T_j and let all other data be induced.
- (3) If e is an external edge of T_j whose internal vertex v is at least trivalent in T_j , then contract e in T_j and then perform a vertex explosion at the identified vertex. This procedure yields a graph $(T_j/e)^{[v]}$ which has components $\{T_a\}$, each of which comes equipped with a map $T_a \rightarrow (T_j/e)^{[v]} \rightarrow T_j/e \rightarrow \Gamma/e$. This specifies a subgraph \tilde{T}_a for each T_a . Let all other data be induced. See Figure 9.

Note that if e is an edge of T_j with two external vertices, then T_j is a segment and thus has leaf length 1 and length 1 so e cannot have length 0.

Definition 2.10. Let Γ be a string diagram and let h be a half-edge in T_j whose source is at least trivalent in T_j . Suppose further that the branch T_h of T_j is prunable. The *pruning* of Γ along h , denoted $\Gamma \setminus h$, has the same underlying pseudometric fatgraph as Γ but has its subgraphs changed as follows. The inclusions of T_h and T^h into T_j induce maps from T_h and T^h into Γ which specify \tilde{T}_h and \tilde{T}^h to replace \tilde{T}_j in $\Gamma \setminus h$. See Figure 10.

It is an elementary lemma to verify that the induced structures described are in fact string diagrams. That is, it is direct to show that the balance conditions for the fatgraph tree subgraphs and the existence, uniqueness, and cycle-free properties of fundamental vertices are preserved in the edge contraction or the pruning.

Remark. Let Γ be a string diagram. For an edge e of length zero, there is a canonical isomorphism from $|\Gamma|$ to $|\Gamma/e|$; similarly, for a half-edge h which determines a prunable branch in some T_j , there is a canonical isomorphism from $|\Gamma|$ to $|\Gamma \setminus h|$.




FIGURE 9. Contraction of an external edge of a tree whose internal vertex is 3-valent.




FIGURE 10. The pruning of a tree T_j in a string diagram along h .

In the case that Γ is an oriented string diagram and e is an internal edge of length zero, we induce an orientation on Γ/e as follows. In cases (1) and (2) in Definition

2.9, an orientation on Γ induces an evident orientation on Γ/e . In case (3), where e is an external edge of a tree T_j whose internal vertex v has valence $|v| \geq 3$, we must be more explicit. Fix a representative ordering of the orientation on Γ so that the tree T_j is first in the ordering on \mathcal{T} . In Γ/e , the tree T_j is replaced by the set of trees $\{T_a\}$; choose an arbitrary order $\{T_a^1, \dots, T_a^{|v|-1}\}$ on this set. The corresponding set of new leaves $\{\ell_a\}$ are added to \mathcal{L} such that $\ell_a^2, \dots, \ell_a^{|v|-1}$ appear in order at the beginning, and the leaf ℓ of e is replaced by the leaf ℓ_a^1 . Notice first that the orientation class of the constructed ordering is invariant of both the representative ordering of Γ and the chosen order on $\{T_a\}$ (because any permutation on the set $\{T_a\}$ induces a corresponding permutation with the same sign on the set $\{\ell_a\}$).

In the case that Γ is an oriented string diagram and h is a half-edge of T_j corresponding to a prunable branch T_h , we induce an orientation on $\Gamma \setminus h$ as follows. Choose a representative ordering of the orientation on Γ so that the tree T_j is first in the ordering on \mathcal{T} . Then replace the tree T_j with the ordered pair (T^h, T_h) . In the ordering on \mathcal{L} , the new leaf $s(h)$ of T_h is added at the beginning.

Remark. The only automorphism of a combinatorial string diagram (fixing the orders on the sets of subgraphs $\{Q_i\}$ and $\{L_i\}$) is the identity. This implies that isomorphic combinatorial string diagrams can be canonically identified. Because of this we can safely be sloppy about whether we are referring to a particular combinatorial string diagram or its isomorphism class. The same is true for oriented combinatorial string diagrams and orientation preserving automorphisms. Our main application of this fact is the canonical identification of the polytopes K for two oriented combinatorial string diagrams in the same isomorphism class.

Definition 2.11. Let \mathcal{SD} be the cell complex built as follows.

There is a proper class of polytopes K , where the underlying oriented combinatorial string diagram ranges over all oriented combinatorial string diagrams. The set of cells of \mathcal{SD} is this proper class under the canonical identification of polytopes for isomorphic oriented combinatorial string diagrams.

The attaching map for K takes the oriented string diagram Γ in a codimension one face of K to the oriented string diagram described below. This is well defined by Proposition 2.12.

- (1) If Γ has an internal edge e of length zero, identify Γ with Γ/e with the induced orientation. Note that this is not a codimension one identification if e is an edge of \tilde{T}_j between an at least trivalent fundamental vertex of \tilde{T}_j and a non-fundamental vertex of \tilde{T}_j .
- (2) If a branch T_h of T_j in Γ is prunable, identify Γ with $\Gamma \setminus h$ with the induced orientation.

Figure 11 shows a picture of the attaching maps.

Proposition 2.12. *The attaching maps in Definition 2.11 are well defined and so \mathcal{SD} is a cell complex.*

Proof. In order to ensure that the attaching maps are well defined, we must check that they agree on codimension two faces. Because each cell K is a polytope, each




FIGURE 11. Attachments of the types listed in Definition 2.11. Compare to Figure 8.

codimension two face arises in precisely two ways as a codimension one face of a codimension one face.

In almost every case, the two codimension one degenerations evidently commute up to a potential difference in orientation. When both codimension one degenerations are pruning degenerations, there are four cases depending on the combinatorics of the trees involved. In every case, it is easy to see that the orientations of the two sequences of attaching maps agree. See Figure 12.

As a representative example, consider case (c) in the figure. Let T be the depicted tree, imagined as some local piece of an oriented string diagram. There is a codimension one degeneration corresponding to pruning T at h ; the orientation $(T, \dots); (\dots)$ induces the orientation $(T^h, T_h, \dots); (s(h), \dots)$. Then pruning T^h at h' induces the orientation

$$((T^h)^{h'}, (T^h)_{h'}, T_h, \dots); (s(h'), s(h), \dots).$$

On the other hand, pruning T at h' induces the orientation $(T^{h'}, T_{h'}, \dots); (s(h'), \dots)$. Further pruning $T_{h'}$ at h requires interchanging $T^{h'}$ and $T_{h'}$ in the representative ordering. Then the induced orientation on the doubly pruned string diagram is

$$((T_{h'})^h, (T_{h'})_h, T^{h'}, \dots); (s(h), s(h'), \dots),$$

where we have swapped $s(h)$ and $s(h')$ in the ordering to compensate for interchanging $T_{h'}$ and $T^{h'}$. Finally, note that $(T^h)^{h'} = T^{h'}$, that $(T^h)_{h'} = (T_{h'})^h$, and that $(T_{h'})_h = T_h$; then the representative orders of the two induced orientations differ by an even permutation on the set of trees and thus agree.

The other cases in Figure 12 are similar and will be omitted.

Only one further case requires comment. If the codimension two face arises by contracting an external edge e of length zero of a tree T_j whose internal vertex v in T_j is trivalent, then the two degenerations are as follows. First recall that branches of T_j correspond to internal half-edges of T_j . At v there are two branches corresponding to the edges other than e . Because e has length zero, both of these




FIGURE 12. The four types of codimension two degenerations where both codimension one degenerations are prunings.

branches are prunable. After performing an identification of type (2) with either one of these branches, the other degeneration becomes a degeneration of type (1), where e is now a length zero edge. The two resulting string diagrams are equivalent. See Figure 13. Given an ordering on a string diagram Γ containing such a tree T_j , the induced orderings for these two degenerations always differ by a transposition on both the tree factor and on the leaf factor, so the induced orientations also agree.

□

Remark. The attaching maps work just as well without an orientation and we can also define a space \mathcal{SD}^u of unoriented string diagrams just as in Definition 2.11. There is a double cover $\mathcal{SD} \rightarrow \mathcal{SD}^u$ given by forgetting the orientation.²

Conjecture. The double cover $\mathcal{SD} \rightarrow \mathcal{SD}^u$ is trivial.

If we knew Conjecture 2.2, we could treat \mathcal{SD}^u as our fundamental space of operations instead of \mathcal{SD} ; as it is, we will need the extra data of the orientation to make choices to build the diffuse intersection class later. See Section 7.

Definition 2.13. Let $\mathcal{SD}(\chi, k, \ell)$ be the subspace of \mathcal{SD} where there are k inputs, ℓ outputs, and where each string diagram has Euler characteristic χ .

Proposition 2.14. The space $\mathcal{SD}(\chi, k, \ell)$ is a finite cell complex.

Proof. None of the attaching maps change the Euler characteristic of the underlying fatgraph or change the number of inputs or outputs. Therefore $\mathcal{SD}(\chi, k, \ell)$

²In the special case where a string diagram has no T_j , then this is *not* a double cover but an isomorphism. See Proposition 8.13 for some discussion of this special case.




FIGURE 13. The induced orientations for a codimension two degeneration that arises by contracting an external edge of a tree T_j whose internal vertex is trivalent.

is a subcomplex of \mathcal{SD} . There are a finite number of connected marked fatgraph isomorphism types with no bivalent vertices and fixed Euler characteristic. Given a marked fatgraph, there are only finitely many ways to give it an oriented combinatorial string diagram structure as in Definition 2.4.

This shows that the space $\mathcal{SD}(\chi, k, \ell)$ is a finite cell complex.

□

Remark. Spaces of fatgraphs like \mathcal{SD} and \mathcal{SD}^u have been used for some time to study not only string topology, but moduli spaces of Riemann surfaces with boundary as well [Str84, Pen87, Har88, Kon92, Igu02, Cos06, God07b]. We will denote the disjoint union of such moduli spaces over all genera and number of incoming and outgoing boundary components by \mathcal{M} . For example, Cohen and Godin defined *Sullivan chord diagrams* and *marked metric chord diagrams*, which they use to define string topology operations [CG04] (we will show in Section 8 that our constructions recover those of Cohen–Godin). Their space of marked metric chord diagrams, which we denote by \mathcal{CG} , includes into \mathcal{M} . For the subspace of genus-zero surfaces with one outgoing boundary component, this inclusion is a homotopy equivalence. Cohen and Godin initially thought that their inclusion could be a homotopy equivalence in general, but this turned out not to be the case [God04]. Tradler and Zeinalian later defined a more general version of marked metric chord diagrams to study algebraic string topology operations [TZ06]. Their space, which we denote by \mathcal{TZ} , is a compactification of the space \mathcal{CG} and in turn, their space includes into a compactification of \mathcal{M} . More specifically, the space \mathcal{TZ} is a deformation retract of Bødigheimer’s harmonic compactification of \mathcal{M} , which we denote by \mathcal{BM} [Böd06, Poi10, EK]. Thus there is a commutative square of inclusions, as in the outside square of the diagram below.

$$\begin{array}{ccc}
\mathcal{CG} & \xleftarrow{\sim} & \mathcal{M} \\
\downarrow & \nearrow \mathcal{SD}^u \quad \sim & \downarrow \\
\mathcal{TZ} & \xleftarrow[q]{\sim} & \mathcal{BM}
\end{array}$$

The space \mathcal{TZ} receives a quotient map from our space \mathcal{SD}^u . We have seen that, for some components with low genus and small numbers of boundary components, the composition of this quotient with the inclusion of \mathcal{TZ} into \mathcal{BM} has a lift up to homotopy to a map from \mathcal{SD}^u to \mathcal{M} , as depicted in the digram. Further, in these examples, this lift is a homotopy equivalence. This constitutes our evidence for Conjecture A in the introduction.

3. STRAIGHTENING A STRING DIAGRAM

Next, we define *straightening* maps from the intersection graph of a string diagram to a product of simplices, one simplex for each component of the intersection graph. Once this is done we will use maps of these simplices into M to define the map \heartsuit and its domain S .

Proposition 3.1. *Given a short-branched tree T with leaf set $L(T)$, let $\Delta_{L(T)}$ be the simplex spanned by $L(T)$. There exists a straightening map str from the pseudometric realization $|T|$ of T to $\Delta_{L(T)}$, which satisfies the following properties.*

- (1) *The map str takes each leaf in $|T|$ to itself in $\Delta_{L(T)}$.*
- (2) *Let e be an internal edge of T or an external edge of T whose internal vertex is bivalent. Assume e has length zero. Let $|T| \rightarrow |T/e|$ be the isomorphism induced by the contraction of the edge e . Then the following diagram commutes.*

$$\begin{array}{ccc}
|T| & \xrightarrow{\text{str}} & \Delta_{L(T)} \\
\downarrow & & \downarrow \cong \\
|T/e| & \xrightarrow{\text{str}} & \Delta_{L(T/e)}
\end{array}$$

- (3) *Let T_h be a prunable branch of T with pollard T^h . Since $L(T^h)$ is a subset of $L(T)$, there is a natural inclusion of $\Delta_{L(T^h)}$ in $\Delta_{L(T)}$. Since every leaf of T_h except $s(h)$ is also a leaf of T , assigning a point in $\Delta_{L(T)}$ to $s(h)$ yields a linear inclusion of $\Delta_{L(T_h)}$ in $\Delta_{L(T)}$. Here, since $s(h)$ is also a vertex of T^h , it has an image point $|s(h)|$ in the pseudometric realization $|T^h|$. Thus we assign to $s(h)$ the image of $|s(h)|$ under the straightening map of T^h .*

Then the following diagram commutes.

$$\begin{array}{ccc}
 |T^h| \sqcup |T_h| & \xrightarrow{\text{str}} & \Delta_{L(T^h)} \sqcup \Delta_{L(T_h)} \\
 \downarrow & & \downarrow \\
 |T| & \xrightarrow{\text{str}} & \Delta_{L(T)}
 \end{array}$$

The proof is technical and is deferred to Appendix A. See Figure 14 for a picture of the straightening map.




FIGURE 14. A short-branched tree T and its straightening in $\Delta_{L(T)}$

Let C be a component of the intersection graph $\widehat{\Gamma}$ of the string diagram Γ with leaf set $L(C)$. Let $\Delta_{L(C)}$ be the simplex generated by leaves of C . Here we describe how to use the straightening map for trees above to extend the inclusion of the leaves of C into $\Delta_{L(C)}$ to a map from $|C|$ to $\Delta_{L(C)}$, even if C is not a tree.

The component C comes equipped with canonical maps from each tree in a subset of the set of trees $\{T_j\}$ discussed in the remark following Definition 2.3. We will give a well-defined map from $|C|$ to $\Delta_{L(C)}$, defined piecewise by first mapping each of these trees $|T_j|$ to $\Delta_{L(T_j)}$ using straightening maps above, and then mapping each such $\Delta_{L(T_j)}$ to $\Delta_{L(C)}$.

The maps $\Delta_{L(T_j)}$ to $\Delta_{L(C)}$ are defined inductively and we require an ordered partition P_1, P_2, \dots, P_m of the set of trees T_j of C . The ordered partition comes from the canonical map from T_j to C and is defined as follows. The tree T_j is in P_n if it is not in any previous subset of the partition and the image in C of the leaves of T_j is contained in the union of the leaves of C and the images in C of all trees in P_1, P_2, \dots, P_{n-1} .

Assume the maps $\Delta_{L(T_j)}$ to $\Delta_{L(C)}$ have been defined for trees in P_1, P_2, \dots, P_{n-1} . Let T_j be in P_n . A vertex of $\Delta_{L(T_j)}$ corresponds to a leaf of T_j , which already has

an image in Δ_{L_C} by the definition of P_n . Extend this linearly to a map from $\Delta_{L(T_j)}$ to Δ_C .

Definition 3.2. The *straightening* of a component C of the intersection graph $\widehat{\Gamma}$ of a string diagram Γ is the map described immediately above from $|C|$ to $\Delta_{L(C)}$. The *straightening* of the string diagram Γ is the disjoint union over components C of the intersection graph $\widehat{\Gamma}$ of the straightening maps $|C|$ to $\Delta_{L(C)}$. We also denote this map by str .

The conditions of Proposition 3.1 may be extended to the straightening maps for string diagrams, which are used in the definition of the map \heartsuit in the next section.

Recall that if a string diagram has an internal edge of length zero, the attaching map of Definition 2.11 contracts the edge. If the string diagram has a tree with a prunable branch, the attaching map breaks the tree into two trees.

Roughly, the next lemma shows that contracting an edge commutes with straightening. For the kind of edges we will consider, there is a canonical bijection between the leaf set of C and the leaf set of C/e .

Lemma 3.3. *Let $|C| \rightarrow |C/e|$ be the map which contracts an edge e which is either a external edge with bivalent internal vertex or an internal edge which is not the image of an external edge of a tree T_j whose internal vertex is at least trivalent in C . Then the following diagram commutes.*

$$\begin{array}{ccc} |C| & \xrightarrow{\text{str}} & \Delta_{L_C} \\ \cong \downarrow & & \downarrow \cong \\ |C/e| & \xrightarrow{\text{str}} & \Delta_{L_{C/e}} \end{array}$$

Proof. The edge e of C is an edge in the image of a unique tree T_j . The straightening maps and simplex inclusion maps for any tree earlier than or incomparable with T_j in the ordered partition of trees in C are unaffected by the edge contraction.

Proposition 3.1 shows that straightening the tree T_j commutes with contracting e . Because all previous straightening maps and simplex inclusions are unchanged, the map $\Delta_{L(T_j)} \rightarrow \Delta_{L(C)}$ is also unchanged. Then, the straightening maps and simplex inclusions are unchanged for trees after T_j in the ordered partition. \square

Next, we consider a string diagram Γ where the tree T_j has a prunable branch T_h . Let C be the component of $\widehat{\Gamma}$ which receives the canonical map from T_j . Under the attaching map in Definition 2.11, Γ is identified with the string diagram $\widehat{\Gamma \setminus h}$ where the tree T_j is broken into two trees T^h and T_h . The component $C \setminus h$ of $\widehat{\Gamma \setminus h}$ which receives the canonical map from T^h and T_h is canonically isomorphic as a pseudometric fatgraph to C .

Lemma 3.4. *Let C and $C \setminus h$ be as above and let $|C| \rightarrow |C \setminus h|$ be the identity map. Then the following diagram commutes.*

$$\begin{array}{ccc}
 |C| & \xrightarrow{\text{str}} & \Delta_{L_C} \\
 \cong \downarrow & & \downarrow \cong \\
 |C \setminus h| & \xrightarrow{\text{str}} & \Delta_{L_{C \setminus h}}
 \end{array}$$

Proof. As in Lemma 3.3, earlier and incomparable trees are not affected a priori. Proposition 3.1 show that T_j and subsequent trees have the same straightening maps and simplex inclusions. \square

4. THE HEART OF THE STRING TOPOLOGY CONSTRUCTION

Let M be a closed Riemannian manifold. Given an oriented string diagram Γ , and a map γ from its inputs into M , we would like to produce a map from its outputs into M . To do this, we would like to extend γ to a map from $|\Gamma|$ to M , and then pull back to the outputs of Γ . The extension of γ to $|\Gamma|$ is our version of a wrong-way map and is the heart of the string topology construction.

Our construction requires certain points in the image of γ be close together in M . In this section, we define

- (1) a space S which realizes this closeness condition,
- (2) a space $\mathcal{SD}(M)$ which captures mapping spaces from oriented string diagrams to M as oriented string diagrams vary in \mathcal{SD} , and
- (3) a map \heartsuit from S to $\mathcal{SD}(M)$ which realizes the extension of γ to $|\Gamma|$.

In what follows, all oriented string diagrams have Euler characteristic χ and k inputs. The number of outputs ℓ is also fixed and to simplify notation, we use \mathcal{SD} for $\mathcal{SD}(\chi, k, \ell)$. All of the constructions in this section would work just as well for unoriented string diagrams, but as we will eventually need to use the orientation, we have chosen to work with it from the beginning.

Let LM be the space of continuous maps from the standard circle into M . Let LM^k be the k -fold Cartesian product of LM with itself. A point in LM^k may be represented as a map γ from k copies of the standard circle into M . In the definition, we will use the canonical identification of the domain of γ with the subspace of the pseudometric realization of an oriented string diagram Γ determined by the subgraph $\sqcup Q_i$. We also specify the notation ι for the map from the leaves of $\hat{\Gamma}$ to $\sqcup |Q_i|$.

The domain of the map \heartsuit will be fibered over \mathcal{SD} and we begin by describing the fiber over a particular string diagram Γ . In fact, we will describe a fiber that depends on a parameter ε and later we will fix the appropriate ε which will make the construction of \heartsuit possible.

Definition 4.1. Let Γ be an oriented string diagram and let γ be in LM^k . Let a and b be two leaves in the same component of the intersection graph $\widehat{\Gamma}$. We say that γ is ε -Lipschitz with respect to a and b if the distance in M between $\gamma(\iota(a))$ and $\gamma(\iota(b))$ is less than ε times the distance in $|\widehat{\Gamma}|$ between a and b . We say that γ is ε -Lipschitz with respect to Γ if it is ε -Lipschitz with respect to all such pairs of leaves.

We call such maps ε -Lipschitz with respect to Γ because we would like to extend γ to a function from $|\Gamma|$ to M whose restriction to $|\widehat{\Gamma}|$ is Lipschitz with Lipschitz constant ε .

This definition depends on the structure of Γ . However, it turns out that contracting edges of length zero and pruning prunable branches do not change the ε -Lipschitz condition.

Lemma 4.2. Let Γ and Γ' be two oriented string diagrams identified by the attaching map of Definition 2.11, and let γ be in LM^k . Then γ is ε -Lipschitz with respect to Γ if and only if γ is ε -Lipschitz with respect to Γ' .

Proof. It suffices to consider the case when Γ' is obtained from Γ by a codimension one degeneration. In fact, there is only one type of degeneration which affects either the leaf set or components of the intersection graph $\widehat{\Gamma}$ or distances in its pseudometric realization $|\widehat{\Gamma}|$. Namely, it is necessary to prove the lemma for the case $\Gamma' = \Gamma/e$ where e is an external edge of $\widehat{\Gamma}$ of length zero. The contraction has the potential to break the component of $\widehat{\Gamma}$ containing the edge e into multiple components.

Let c be the leaf of e and let v be its internal vertex. Let a and b be leaves of $\widehat{\Gamma}$ which are in the same component as c .

Let d denote the distance functions in M , in $|\widehat{\Gamma}|$, and in $|\widehat{\Gamma}'|$. First assume γ is ε -Lipschitz with respect to Γ . Thus for $\widehat{\Gamma}$ we have

$$\begin{aligned} d(\gamma(\iota(a)), \gamma(\iota(b))) &< \varepsilon d(a, b), \\ d(\gamma(\iota(b)), \gamma(\iota(c))) &< \varepsilon d(b, c), \text{ and} \\ d(\gamma(\iota(a)), \gamma(\iota(c))) &< \varepsilon d(a, c). \end{aligned}$$

The contraction of e identifies the leaves a and b of $\widehat{\Gamma}$ with leaves a' and b' of $\widehat{\Gamma}'$. The leaf c may be identified with multiple leaves of $\widehat{\Gamma}'$; exactly one such c' is in the same component as a' and another such c'' is in the same component as b' .

Recall that the distances between vertices in the $|\widehat{\Gamma}|$ are induced by lengths of edges in $\widehat{\Gamma}$ and likewise for $\widehat{\Gamma}'$. Because lengths in Γ' are induced by those in Γ , for $\widehat{\Gamma}'$, we have

$$\begin{aligned} d(\gamma(\iota(b')), \gamma(\iota(c''))) &< \varepsilon d(b', c'') \text{ and} \\ d(\gamma(\iota(a')), \gamma(\iota(c'))) &< \varepsilon d(a', c'). \end{aligned}$$

This shows that if γ is ε -Lipschitz with respect to Γ then γ is also ε -Lipschitz with respect to Γ' .

To see the converse, we assume the previous two inequalities, which yield

$$\begin{aligned} d(\gamma(\iota(b)), \gamma(\iota(c))) &< \varepsilon d(b, c) \text{ and} \\ d(\gamma(\iota(a)), \gamma(\iota(c))) &< \varepsilon d(a, c). \end{aligned}$$

We must verify $d(\gamma(\iota(a)), \gamma(\iota(b))) < \varepsilon d(a, b)$ as well.

If the distance-minimizing path between a and b in $|\widehat{\Gamma}|$ does not pass through the vertex v , then contracting e does not affect the path and because γ is ε -Lipschitz with respect to Γ' , we have $d(\gamma(\iota(a)), \gamma(\iota(b))) < \varepsilon d(a, b)$.

If the distance-minimizing path between a and b in $|\widehat{\Gamma}|$ does pass through the vertex v , then $d(a, b)$ is equal to $d(a, c) + d(b, c)$ since the length of e is zero. In this case, we have that

$$\begin{aligned} d(\gamma(\iota(a)), \gamma(\iota(b))) &\leq d(\gamma(\iota(a)), \gamma(\iota(c))) + d(\gamma(\iota(b)), \gamma(\iota(c))) \\ &< \varepsilon d(a, c) + \varepsilon d(b, c) \\ &= \varepsilon d(a, b). \end{aligned}$$

This concludes the proof. \square

The construction of the map \heartsuit will require that ε be small enough relative to the geometry of our manifold M . In particular, the map \heartsuit is defined using a composition of the straightening maps from Section 3 and maps from the standard simplex into M , which are built using the technique of Riemannian centers of mass³ pioneered by Grove and Karcher [GK73]. These techniques require strongly convex balls in M . In order to guarantee strongly convex balls, we use the bounds in Definition 4.3. The presentation of this material mainly follows [San12]; see [Afs11] for a concise historical review.

Definition 4.3. Let M be a Riemannian manifold. The *pre-convexity radius* of M , denoted r , is

$$r := \frac{1}{2} \min \left\{ \text{inj}(M), \frac{\pi}{\sqrt{\Delta}} \right\}$$

where $\text{inj}(M)$ is the injectivity radius of M and Δ is the supremum of sectional curvatures of M . If $\Delta \leq 0$, we interpret $\frac{\pi}{\sqrt{\Delta}}$ as ∞ .

Theorem 4.4. Let M be a compact Riemannian manifold with pre-convexity radius r . Fix an r -ball B in M . Let F be a finite set and let Δ_F be the simplex spanned by F . Then there exists a smooth map $\Upsilon_F : B^F \times \Delta_F \rightarrow B$ called the simplicial geodesic interpolation such that

- (1) Υ_F , restricted to $B^F \times F$, is evaluation: $(f, x) \mapsto f(x)$,
- (2) simplicial geodesic interpolation, restricted to a face of the simplex, is simplicial geodesic interpolation for that face, and
- (3) fixing a configuration f in M^F , the map $\Upsilon(f, \cdot)$ viewed as a map from Δ_F to M does not depend on the choice of B containing $f(F)$.

Remark. This theorem summarizes Definition 2.2, Theorem 2.1, and Corollary 2.2 of [San12]. There the reliance on injectivity radius is erroneously omitted. The

³Also called Karcher means. See [Kar14].

correct choice of r and independence of B are given in [Afs11, Section 1, Theorem 2.1] with some illuminating discussion.

The values of ε that make the construction of \heartsuit possible are $\varepsilon = \frac{r}{|\chi|}$ and $\varepsilon = \frac{r}{2|\chi|}$. We use the larger value of ε to define the space S , the domain of \heartsuit , and we use the smaller value to identify a subspace s of S which we will use later for excision.

Definition 4.5. The *diffuse intersection locus* is the subspace S of $\mathcal{SD} \times LM^k$ consisting of pairs (Γ, γ) where γ is $\frac{r}{|\chi|}$ -Lipschitz with respect to Γ .

We also identify the subspace s of S consisting of pairs (Γ, γ) where γ is $\frac{r}{2|\chi|}$ -Lipschitz with respect to Γ .

Remark. The closure of s is contained in S .

The map \heartsuit will have codomain $\mathcal{SD}(M)$, which is a universal space over \mathcal{SD} .

Definition 4.6. The space $\mathcal{SD}(M)$ consists of pairs (Γ, Θ) where Γ is an oriented string diagram and Θ is a map from $|\Gamma|$ to M .

Remark. The set $\mathcal{SD}(M)$ comes equipped with a forgetful map to the space \mathcal{SD} , $(\Gamma, \Theta) \mapsto \Gamma$. The fiber over Γ is $\text{Maps}(|\Gamma|, M)$; the fiber over a point in a small neighborhood of Γ comes with a canonical map to the fiber over Γ . We use standard point-set techniques to build a basis generating a topology on $\mathcal{SD}(M)$ from the topologies on \mathcal{SD} and on fibers. The notation $\mathcal{SD}(M)$ refers to this topological space.

We are ready to define the map $\heartsuit : S \rightarrow \mathcal{SD}(M)$. For a pair (Γ, γ) , the map \heartsuit will give us the pair $(\Gamma, \Theta(\Gamma, \gamma))$ where $\Theta(\Gamma, \gamma)$ is a map from $|\Gamma|$ to M extending γ . To define $\Theta(\Gamma, \gamma)$, we need only specify its behavior on the pseudometric realization of the intersection graph $|\widehat{\Gamma}|$ and ensure it agrees with γ on its leaves.

Definition 4.7. The restriction of $\Theta(\Gamma, \gamma)$ to $|\widehat{\Gamma}|$ is defined as follows. Let C be a component of $\widehat{\Gamma}$ with leaf set $L(C)$. First we straighten C to get a map from $|C|$ to $\Delta_{L(C)}$. Next, we use γ to specify a map from the vertices of $\Delta_{L(C)}$ into M and apply Υ_C to map $\Delta_{L(C)}$ into M .

Formally, for x a point in $|C|$, we have

$$\Theta(\Gamma, \gamma)(x) = \Upsilon_{L(C)}(\gamma \circ \iota, \text{str}(x)).$$

The attaching map of Definition 2.11 which identifies oriented string diagrams Γ and Γ' induces a canonical identification of their pseudometric realizations $|\Gamma|$ and $|\Gamma'|$. We would like to show that if γ is ε -Lipschitz with respect to Γ and Γ' , then $\Theta(\Gamma, \gamma)$ and $\Theta(\Gamma', \gamma)$ are equal. We will show this for particular pairs of oriented string diagrams Γ and Γ' identified under the attaching map; transitivity will imply that it is true for all such pairs.

More specifically, let Γ be an oriented string diagram in the boundary of a cell K which is identified with the oriented string diagram Γ' in the interior of a cell K' . The choice of a sequence of codimension one attaching maps identifying Γ with Γ' yields a canonical isomorphism from $|\Gamma|$ to $|\Gamma'|$. This isomorphism is independent of the choice of the sequence.

Proposition 4.8. *Let Γ be in the boundary of K and let Γ' be in K' as above. Assume that γ is ε -Lipschitz with respect to Γ and Γ' . Then the maps $\Theta(\Gamma, \gamma)$ and $\Theta(\Gamma', \gamma)$, from $|\Gamma| \cong |\Gamma'|$ to M , are equal.*

Proof. The oriented string diagram Γ' is obtained from the oriented string diagram Γ by contracting internal edges of length zero and pruning prunable branches. There are three different types of internal edges of length zero, so in total there are four types of degenerations to consider.

- (1) The intersection graph $\widehat{\Gamma}$ has an edge of length zero which is either an internal edge or a external edge whose internal vertex is bivalent.
- (2) The subgraph $\sqcup Q_i$ of Γ has an internal edge of length zero.
- (3) The intersection graph $\widehat{\Gamma}$ has an external edge of length zero whose internal vertex is at least trivalent.
- (4) The oriented string diagram Γ has a tree with a prunable branch.

Lemmas 3.3 and 3.4 show that the straightening of Γ commutes with contracting length zero edges of $\widehat{\Gamma}$ of types listed in case (1) and pruning branches as in case (4). In these cases, the leaf set of each component C of $\widehat{\Gamma}$ is unchanged, so the map $\iota : L(\widehat{\Gamma}) \rightarrow \sqcup |Q_i|$ is also unchanged. This shows that the maps $\Theta(\Gamma, \gamma)$ and $\Theta(\Gamma', \gamma)$ are equal for cases (1) and (4).

There is a special case of case (1) that appears to be missing: where the zero-length edge is a external edge of some tree T_j whose internal vertex is at least trivalent in T_j . In fact, in this case there is always a prunable branch T_h of T_j such that the source of h is this vertex. Thus, after repeated applications of Lemma 3.4, this case is then accounted for by Lemma 3.3.

In case (2), the straightening maps of Γ and Γ' , from $|\widehat{\Gamma}| \cong |\widehat{\Gamma}'|$ to the disjoint unions $\sqcup \Delta_{L(C)} \cong \sqcup \Delta_{L(C')}$, coincide. This shows that the maps $\Theta(\Gamma, \gamma)$ and $\Theta(\Gamma', \gamma)$ are equal for case (2).

Case (3) will take the most care because contracting an external edge e of the intersection graph of Γ whose internal vertex v is at least trivalent changes the leaf set of the intersection graph and may even break $\widehat{\Gamma}$ into more components.

In this case, contraction of the edge e induces a natural map φ from $\widehat{\Gamma}'$ to $\widehat{\Gamma}$. The map φ identifies corresponding edges of $\widehat{\Gamma}'$ with edges of $\widehat{\Gamma}$ and identifies corresponding leaves of $\widehat{\Gamma}'$ with leaves of $\widehat{\Gamma}$ or the vertex v . The number of leaves of $\widehat{\Gamma}'$ sent to v by φ is equal to one less than the valence of v ; otherwise φ is injective. In Figure 9, $\widehat{\Gamma}$ is equal to \widetilde{T}_j and $\widehat{\Gamma}'$ is equal to the union of \widetilde{T}_a and $\widetilde{T}_{a'}$. In this example, the map φ sends one leaf from each of these subgraphs to the vertex v .

To complete the proof, we show that contracting the edge commutes first with straightening and then with the Riemannian center of mass map in the appropriate sense.

Let $\{C\}$ and $\{C'\}$ be the sets of components of $\widehat{\Gamma}$ and $\widehat{\Gamma}'$ respectively.

There is an induced natural map Φ from $\sqcup\Delta_{L(C')}$ to $\sqcup\Delta_{L(C)}$ so that the diagram commutes.

$$\begin{array}{ccc} |\widehat{\Gamma}'| & \xrightarrow{\text{str}} & \sqcup\Delta_{L(C')} \\ \downarrow |\varphi| & & \downarrow \Phi \\ |\widehat{\Gamma}| & \xrightarrow{\text{str}} & \sqcup\Delta_{L(C)} \end{array}$$

We will define Φ on the vertices of $\sqcup\Delta_{L(C')}$, that is, the leaves of $\widehat{\Gamma}'$, and extend linearly in each simplex factor. Any leaf of $\widehat{\Gamma}'$ is taken by φ either to a leaf of $\widehat{\Gamma}$ or to the internal vertex of e . The map Φ takes vertices in the first case to their image under φ , and takes vertices in the second case to the leaf vertex of e .

Since e has length zero, the straightening map for Γ takes both vertices of e to the same point in $\sqcup\Delta_{L(C)}$, namely, to the leaf vertex of e . This shows that contracting e commutes with straightening.

Finally, since Φ takes vertices to vertices we have the following commutative diagram.

$$\begin{array}{ccccc} \sqcup L(C') & & & & \\ \downarrow \Phi & \searrow \sqcup\iota' & & \searrow \gamma & \\ & & \sqcup S^1 & \longrightarrow & M \\ \downarrow \Phi & \nearrow \sqcup\iota & & & \\ \sqcup L(C) & & & & \end{array}$$

This shows the maps $\Theta(\Gamma, \gamma)$ and $\Theta(\Gamma', \gamma)$ are equal for case (3).

□

Definition 4.9. The function $\heartsuit : S \rightarrow \mathcal{SD}(M)$ is defined as $\heartsuit(\Gamma, \gamma) = (\Gamma, \Theta(\Gamma, \gamma))$.

Remark. The function \heartsuit is well-defined by Proposition 4.8. The map \heartsuit is continuous because of the following. First, consider the the image of $K \times LM^k$ in $\mathcal{SD} \times LM^k$ under the characteristic map. The restriction of \heartsuit to the intersection of S with this subspace depends continuously on parameters in K . Then, because this intersection is closed in S , the map \heartsuit is continuous by the gluing lemma.

5. THE PUSH-PULL MAP AND THE STRING TOPOLOGY CONSTRUCTION

All but one of the ingredients for the string topology construction are now in place. In this section we fix an arbitrary singular cochain W in $C^*(S, S - s)$ and define a map called the *push-pull map* for W . We will eventually restrict to the subset of cocycles representing a particular homology class in $H^{|\chi|d}(S, S - s)$, called the *diffuse intersection class*. The push-pull map for a cocycle representing the diffuse intersection class will be called a *string topology construction*. The definition of the diffuse intersection class is rather involved, and its definition is the goal of Sections 6 and 7. Again, everything in this section would work fine if we replaced \mathcal{SD} with its unoriented version \mathcal{SD}^u ; we will only use the orientation to build the

diffuse intersection class. We work with coefficients in R , an arbitrary commutative ring with identity; we suppress the notation throughout.

Let M^d be a closed, R -oriented d -dimensional Riemannian manifold. Then for each χ , k , and ℓ , the push-pull map for a degree $|W|$ cochain W in $C^*(S, S - s)$ is a degree $-|W|$ map

$$\mathcal{ST}_W : C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k} \rightarrow C_*(LM)^{\otimes \ell},$$

which we express as a composition of two maps, given in Definitions 5.1 and 5.2 below.

The first map in this composition is our version of a wrong-way map given by a Pontryagin–Thom construction.

Definition 5.1. Let

$$(\rho_{in})! : C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k} \rightarrow C_{*-|W|}(\mathcal{SD}(M))$$

be the following composition of maps:

- (1) first, $C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k} \rightarrow C_*(\mathcal{SD} \times LM^k)$, an Eilenberg–Zilber map for singular chains,
- (2) next, $C_*(\mathcal{SD} \times LM^k) \rightarrow C_*(\mathcal{SD} \times LM^k, \mathcal{SD} \times LM^k - s)$, the usual quotient map from absolute to relative chains,
- (3) next, $C_*(\mathcal{SD} \times LM^k, \mathcal{SD} \times LM^k - s) \rightarrow C_*(S, S - s)$ a chain homotopy inverse to the map induced on chains by the inclusion of spaces $(S, S - s) \hookrightarrow (\mathcal{SD} \times LM^k, \mathcal{SD} \times LM^k - s)$
- (4) next, $C_*(S, S - s) \rightarrow C_{*-|W|}(S)$, given by the cap product with the relative cochain W , and
- (5) finally, $\heartsuit_* : C_{*-|W|}(S) \rightarrow C_{*-|W|}(\mathcal{SD}(M))$, the map induced on singular chains by the map of spaces $\heartsuit : S \rightarrow \mathcal{SD}(M)$.

For concreteness, we choose the so-called shuffle map [Dol95, VI.12.26.2] as our Eilenberg–Zilber map and use the explicit formula implementing iterated barycentric subdivision in [Hat02, Theorem 2.20] as our chain homotopy inverse to the map induced on chains by inclusion.

Remark. All of the maps involved in this definition are chain maps except potentially the cap product with W , which is only a chain map if W is a relative cocycle.

The second map in the composition defining \mathcal{ST}_W is also itself a composition, albeit a simpler one. First recall that for a string diagram Γ with marked output boundary cycle C we have the map $\bar{\partial}_C$ from the standard circle to $|\Gamma|$ which transverses the oriented edges of C in “reverse order.”

Definition 5.2. Let

$$(\rho_{out})_* : C_*(\mathcal{SD}(M)) \rightarrow C_*(LM)^{\otimes \ell}$$

be the following composition of chain maps:

- (1) first, $C_*(\mathcal{SD}(M)) \rightarrow C_*(LM^\ell)$, induced by the map of spaces ρ_{out} from $\mathcal{SD}(M)$ to LM^ℓ which takes $(\Gamma, \Theta : |\Gamma| \rightarrow M)$ to the pullback of Θ to its outputs along the disjoint union of $\bar{\partial}_C$ over all output boundary cycles of Γ , and

(2) next, $C_*(LM^\ell) \rightarrow C_*(LM)^{\otimes \ell}$, an Eilenberg–Zilber map.

For concreteness we choose the Alexander–Whitney map as our Eilenberg–Zilber map in this direction.

We come to the main definition of the paper.

Definition 5.3. Let W be a cochain in $C^*(S, S - s)$. The *push-pull map for W* \mathcal{ST}_W is the degree $-|W|$ map given by the composition of the above maps:

$$\mathcal{ST}_W := (\rho_{out})_* \circ (\rho_{in})! : C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k} \rightarrow C_{*-|W|}(LM)^{\otimes \ell}.$$

Proposition 5.4. *If W is a cocycle in $C^*(S, S - s)$ then \mathcal{ST}_W is a chain map. If W and W' are cohomologous cocycles in $C^*(S, S - s)$ then $\mathcal{ST}_{W'}$ and \mathcal{ST}_W are chain homotopic chain maps.*

Proof. If W is a cocycle then \mathcal{ST}_W is a composition of chain maps.

Let \widetilde{W} be a cochain with coboundary $W - W'$. Then $\mathcal{ST}_{\widetilde{W}}$ is a chain homotopy between \mathcal{ST}_W and $\mathcal{ST}_{W'}$. \square

Remark. This proof shows that the string topology construction $\mathcal{ST}(\cdot)$ is a chain map from $C^{-*}(S, S - s)$ to $\text{Hom}(C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k}, C_*(LM)^{\otimes \ell})$, the homomorphism complex.

Sections 6 and 7 are devoted to picking out a particular cohomology class Ω in $H^{|\chi|^d}(S, S - s)$ called the *diffuse intersection class*.

Definition 5.5. The *string topology construction for W* is the push-pull map \mathcal{ST}_W where W is a cocycle representing the diffuse intersection class Ω .

6. PATCHING COHOMOLOGY CLASSES OVER THE SPACE OF ORIENTED STRING DIAGRAMS

Our next goal is to define a cohomology class in $H^{|\chi|^d}(S, S - s)$. This relative cohomology class on S will be defined as the pullback of a cohomology class on a *stratified pair* $\mathcal{N}(\mathcal{SD})$ under an evaluation map from S to $\mathcal{N}(\mathcal{SD})$.

6.1. The stratified pair. In this section we define the pair $\mathcal{N}(\mathcal{SD})$ as a colimit of a diagram in pairs of spaces. Just as for the space S , these pairs of spaces will involve some closeness condition in M . We will again employ the preconvexity radius r (see Definition 4.3).

Up until this point, it has not been important for us to distinguish between the cells of \mathcal{SD} , viewed as abstract polytopes, and the cells of \mathcal{SD} , viewed as subspaces of \mathcal{SD} . In this section, by the *cell* K , we mean the abstract polytope, which comes equipped with a characteristic map $\text{Ch}_K : K \rightarrow \mathcal{SD}$. Additionally, by a *face* of K , we mean a face \widehat{K}' of the abstract polytope K . The face \widehat{K}' comes equipped with a canonical identification $\alpha : \widehat{K}' \xrightarrow{\cong} K'$, with a particular cell K' of \mathcal{SD} . By a *degeneration* of K , we mean such a cell K' , equipped with the inclusion $K' \xrightarrow{\cong} \widehat{K}' \hookrightarrow K$, whose image is \widehat{K}' .

Definition 6.1. Fix an oriented combinatorial string diagram and the corresponding cell K of \mathcal{SD} . Let \mathcal{T} denote the set of trees $\{T_j\}$ of the combinatorial string diagram and let \mathcal{L} denote the disjoint union of the leaf sets $L(T_j)$. Both of these sets depend on K . Let N_K be the subspace of $M^{\mathcal{L}}$ consisting of maps f from \mathcal{L} to M such that for all trees T_j in \mathcal{T} , the image of the restriction of f to $L(T_j)$ lies in an r -ball in M .

Next we want to discuss the combinatorics of N_K and how it varies as we move around in \mathcal{SD} . To this end, we consider codimension one degenerations K' of a cell K . Recall that there are two types of codimension one degenerations, one given by contracting edges and one given by pruning branches. We call these *contraction* and *pruning* degenerations respectively.

For a codimension one degeneration K' of K , there is canonical map $\mathcal{T}' \rightarrow \mathcal{T}$ between the sets of trees of the corresponding oriented combinatorial string diagrams. This map is bijective for a contraction degeneration. In a pruning degeneration, the map is surjective and is one-to-one outside of T^h and T_h ; it sends both T^h and T_h in \mathcal{T}' to the tree T_j in \mathcal{T} .

This canonical map of sets of trees does not, in general, induce a map of sets of leaves. However, it does give rise to a map in the other direction $\xi : \mathcal{L} \rightarrow \mathcal{L}'$ where a leaf of a tree in \mathcal{T} is sent to the corresponding leaf of a tree in \mathcal{T}' . This map is also bijective for a contraction degeneration. In a pruning degeneration, ξ is injective but misses exactly one leaf w of one tree T_h in \mathcal{T}' , namely, the source of h in T_h .

Note that if K'_1 and K'_2 are codimension one degenerations of the cell K , and K'' is a common codimension one degeneration of K'_1 and K'_2 , then the compositions $\mathcal{L} \rightarrow \mathcal{L}'_1 \rightarrow \mathcal{L}''$ and $\mathcal{L} \rightarrow \mathcal{L}'_2 \rightarrow \mathcal{L}''$ need not agree; the following diagram need not commute. See Figure 15. This possible discrepancy will be shown not to matter for our purposes.

$$\begin{array}{ccc} \mathcal{L} & \longrightarrow & \mathcal{L}'_1 \\ \downarrow & & \downarrow \\ \mathcal{L}'_2 & \longrightarrow & \mathcal{L}'' \end{array}$$

We will use ξ and the Riemmanian center of mass map Υ to build a kind of attaching map for building blocks of the form $K \times N_K$. This attaching map, roughly, extends a map $\mathcal{L} \rightarrow M$ to a map $\mathcal{L}' \rightarrow M$. We begin by defining a map $\nabla^{\widehat{K}'}$ from $\widehat{K}' \times N_K$ to $K' \times M^{\mathcal{L}'}$ and in Lemma 6.4 show that its image lies in $N_{K'}$.

For a pair (Γ, f) in $\widehat{K}' \times N_K$, the map $\nabla^{\widehat{K}'}$ will give us the pair $(\alpha(\Gamma), \Psi(\Gamma, f))$ where $\Psi(\Gamma, f)$ is a map from $\mathcal{L}_{K'}$ to M extending $f\xi^{-1}$. For a contraction degeneration, this suffices to define $\Psi(\Gamma, f)$. For a pruning degeneration, we need only specify the value of $\Psi(\Gamma, f)$ on the leaf w that is not in the image of ξ .

Definition 6.2. Let (Γ, f) be in $\widehat{K}' \times N_K$, where K' is a pruning degeneration of K . The value of $\Psi(\Gamma, f)$ on w is defined as follows. Since w is a leaf of T_h , it is also a vertex of T_j . In the following, str denotes the straightening map $|T_j| \rightarrow \Delta_{L(T_j)}$ of Proposition 3.1.




FIGURE 15. This figure shows the noncommutativity of composition of ξ maps for two codimension two degenerations. Length-zero edges are labeled. The leaf for which the values of the two compositions are different is highlighted.

Let ι be the inclusion of $L(T_j)$ into \mathcal{L} . Then formally,

$$\Psi(\Gamma, f)(w) = \Upsilon_{L(T_j)}(f \circ \iota, \text{str}(w)).$$

Definition 6.3. The map $\nabla^{\widehat{K}'} : \widehat{K}' \times N_K \rightarrow K' \times N_{K'}$ is defined as

$$\nabla^{\widehat{K}'}(\Gamma, f) = (\alpha(\Gamma), \Psi(\Gamma, f)).$$

Compare these definitions to Definitions 4.7 and 4.9. In both cases, we use the composition of straightening and the Riemannian center of mass maps. As a result, the maps \heartsuit and ∇ have good compatibility, which will be exploited later.

When K and K' are clear from context we will suppress them and use ∇ to refer to $\nabla^{\widehat{K}'}$. This is the beginning of a relentless campaign of abuse of notation where any map with a passing resemblance to any $\nabla^{\widehat{K}'}$ will be referred to with the notation ∇ .

Lemma 6.4. *The image of ∇ is contained in $K' \times N_{K'}$.*

Proof. For a contraction degeneration, the map ∇ is a homeomorphism to $K' \times N_{K'}$. This is because the condition of being in N_K and in $N_{K'}$ are the same. For a pruning degeneration, the proof is also straightforward. By assumption, the leaves of T_j are sent into an r -ball in M . Therefore, the leaves of T^h are sent into the same r -ball and $\Upsilon_{L(T_j)}(f \circ \iota, \text{str}(w))$ lies in the same r -ball, so the leaves of T_h also lie into the same r -ball. \square

Remark. The map $\nabla : \widehat{K}' \times N_K \rightarrow K' \times N_{K'}$ is injective.




FIGURE 16. The map f' sends leaves of T_j into a $\delta\|T_j\|$ -ball.

We use the notation $\|T\|$ to refer to the total length of the tree T .

Definition 6.5. Let n_K be the subspace of $M^{\mathcal{L}}$ consisting of maps f from \mathcal{L} to M such that the image of the restriction of f to the set of leaves $L(T_j)$ of any tree T_j lies in a $\frac{\|T_j\|}{|\chi|^2} \frac{r}{4}$ ball in M .

Remark. Since $\|T_j\|$ is less than or equal to $|\chi|$, the closure of n_K is a subspace of N_K .

Lemma 6.6. *The map ∇ sends $\widehat{K'} \times (N_K - n_K)$ into $K' \times (N_{K'} - n_{K'})$.*

Proof. We show that the intersection of $(K' \times n_{K'})$ with the image of ∇ is contained in $\nabla(\widehat{K'} \times n_K)$. For a contraction degeneration, the proof is trivial so we consider only pruning degenerations. For notational convenience, let $\delta = \frac{r}{4|\chi|^2}$.

Assume (Γ', f') is in the intersection of $(K' \times n_{K'})$ with the image of ∇ . This means that f' sends the leaves $L(T^h)$ into a $\delta\|T^h\|$ -ball B^h in M and f' sends the leaves $L(T_h)$ into a $\delta\|T_h\|$ -ball B_h in M . Since (Γ', f') is in the image of ∇ , the map f' sends w to $\Upsilon_{L(T_j)}(f \circ \iota, \text{str}(w))$. This means that f' sends w into B^h .

Therefore the union of B^h and B_h is contained in a $\delta(\|T^h\| + \|T_h\|)$ -ball in M . See Figure 16.

But $L(T_j)$ is contained in the union of $L(T^h)$ and $L(T_h)$ and $\|T_j\| = \|T^h\| + \|T_h\|$. Therefore f' sends the leaves of T_j into a $\delta\|T_j\|$ -ball in M . This shows that (Γ', f') is in $\nabla(\widehat{K'} \times n_K)$. \square

Corollary 6.7. *The map ∇ is a map of pairs of spaces*

$$(\widehat{K'} \times N_K, \widehat{K'} \times (N_K - n_K)) \rightarrow (K' \times N_{K'}, K' \times (N_{K'} - n_{K'})).$$

We will need to refer to these types of pairs of spaces a number of times below. To simplify notation, for J a subspace of a cell K of \mathcal{SD} , we denote the pair

$(J \times N_K, J \times (N_K - n_K))$ by $\mathcal{N}(J; K)$ and refer to the pair as *the K -product over J* . In this notation, ∇ is a map of pairs

$$\mathcal{N}(\widehat{K'}; K) \rightarrow \mathcal{N}(K'; K').$$

We arrive at the first main definition of this section, that of our stratified space $\mathcal{N}(\mathcal{SD})$, which is assembled from all such K -products over the cells J and the ∇ maps.

Definition 6.8. Let $\mathcal{N}(\mathcal{SD})$ be the coequalizer of the following diagram

$$\coprod_{K', K} \mathcal{N}(\widehat{K'}; K) \rightrightarrows \coprod_K \mathcal{N}(K; K)$$

Here the first disjoint union is over pairs of cells of \mathcal{SD} where K' is a codimension one degeneration of K , and the second disjoint union is over cells of \mathcal{SD} . The two maps are

- the maps $\mathcal{N}(\widehat{K'}; K) \hookrightarrow \mathcal{N}(K; K)$ induced by the inclusions of faces $\widehat{K'} \hookrightarrow K$ and
- the maps $\nabla : \mathcal{N}(\widehat{K'}; K) \rightarrow \mathcal{N}(K'; K')$.

The projection maps $K \times N_K \rightarrow K$ and $K \times (N_K - n_K) \rightarrow K$ induce a well-defined projection map of pairs of spaces from $\mathcal{N}(\mathcal{SD}) \rightarrow (\mathcal{SD}, \mathcal{SD})$. Given an inclusion $V \hookrightarrow \mathcal{SD}$, the *space over V* , $\mathcal{N}(V)$ is the preimage of (V, V) in $\mathcal{N}(\mathcal{SD})$ under the projection map. (Because the characteristic map $\text{Ch}_K : K \rightarrow \mathcal{SD}$ is typically not injective on ∂K , we avoid the notation $\mathcal{N}(K)$ and use $\mathcal{N}(\text{Ch}_K(K))$ instead.)

Informally, we can think of $\mathcal{N}(\mathcal{SD})$ as being built inductively. If K is an m -cell and we have already built the space over the $(m-1)$ -skeleton of \mathcal{SD} , then we glue $\mathcal{N}(K; K)$ in using the various ∇ maps as some sort of attaching maps.

A priori, this metaphor is not justified because the attaching could be poorly behaved. This is because of the fact that if K'_1 and K'_2 are codimension one degenerations of K which share a common codimension one degeneration K'' , the compositions of maps ξ of leaves $\mathcal{L} \rightarrow \mathcal{L}'_1 \rightarrow \mathcal{L}''$ and $\mathcal{L} \rightarrow \mathcal{L}'_2 \rightarrow \mathcal{L}''$ need not agree.

As a consequence, it might be possible that gluing in via the “attaching maps” could disrupt the topology of the space over the $(m-1)$ -skeleton of \mathcal{SD} .

Fortunately, this does not occur. In the next lemma, we show that the compositions of the corresponding ∇ maps do agree.

Lemma 6.9. *Let K'_1 and K'_2 be codimension one degenerations of K which share a common codimension one degeneration K'' . Then the two compositions of maps of pairs from the K -product over $\widehat{K''}$ to the K'' -product over K'' —one factoring through the K'_1 -product over $\widehat{K''}$ and the other factoring through the K'_2 -product over $\widehat{K''}$ —agree. Specifically, the following diagram commutes, where maps are understood to refer to their appropriate restrictions:*

$$\begin{array}{ccc}
\mathcal{N}(\widehat{K''}; K) & \xrightarrow{\nabla^{\widehat{K'_1}}} & \mathcal{N}(\widehat{K''}; K'_1) \\
\downarrow \nabla^{\widehat{K'_2}} & & \downarrow \nabla^{\widehat{K''}} \\
\mathcal{N}(\widehat{K''}; K'_2) & \xrightarrow{\nabla^{\widehat{K''}}} & \mathcal{N}(K''; K'')
\end{array}$$

Proof. Contracting a zero length edge of Γ may be a degeneration of codimension greater than one. This phenomenon occurs for an external edge of a tree T_j whose internal vertex is at least trivalent. For the purposes of this proof we call such edges *non-contractible* in Γ .

Consider the codimension two degeneration $K \rightarrow K''$. Pruning degenerations increase the total number of trees in a representative oriented string diagram and all other codimension one degenerations preserve the number of trees. Therefore the number of pruning degenerations in any decomposition $K \rightarrow K' \rightarrow K''$ is constant. We divide into cases.

First, if there are no pruning degenerations, then the statement is trivially true because everything commutes and every map is an isomorphism.

Next, consider a codimension two degeneration sequence $K \rightarrow K'_1 \rightarrow K''$ such that exactly one of the two codimension one degenerations is a pruning degeneration. The other codimension one degeneration must be a contraction degeneration along an edge e of Γ or Γ'_1 .

If e is contractible in Γ , then the unique other degeneration sequence $K \rightarrow K'_2 \rightarrow K''$ giving rise to the same codimension two degeneration consists of the same two codimension one degenerations applied in the opposite order. Contracting zero-length edges does not change the combinatorics of the map ξ on leaf sets, nor do these actions affect the straightening or Riemannian center of mass maps, which give the missing component of $M^{\mathcal{L}''}$.

There is a case involving contracting an edge which is non-contractible in Γ (this is the most involved case). That is, after performing a pruning degeneration on Γ to yield Γ' , there may be an edge which is contractible in Γ' but not in Γ . In this case, the other degeneration sequence involves performing a different pruning degeneration followed by a contraction degeneration. Overall, this codimension two degeneration is realized by the contraction of an external edge of a tree T_j whose non-leaf vertex is trivalent. Recall Figure 15.

In this case, the maps on leaves $\mathcal{L} \rightarrow \mathcal{L}''$ do not commute. There are two leaves v_1 and v_2 of two trees in Γ'' which are identified in Γ'' . The leaves v_1 and v_2 correspond to a vertex v in $\alpha(\Gamma'')$, the oriented string diagram in K .

Examining the compositions $\nabla^{\widehat{K''}} \circ \nabla^{\widehat{K'_2}}$ and $\nabla^{\widehat{K''}} \circ \nabla^{\widehat{K'_1}}$ reveals that the two maps differ only in how they deal with v_1 and v_2 . The first composition identifies v_1 with v and uses the straightening map to decide what to do with v_2 , while the second composition does the reverse. In both cases the second degeneration corresponds to contracting the edge between v_1 and v_2 . By Lemma A.11 this means that the




FIGURE 17. These are examples of the three types of codimension two degenerations arising from distinct pruning degenerations. The pruning edges may be part of the same tree in Γ'' (case a) or may be part of distinct trees. If they are distinct trees, they may share a common pollard (case b) or one may be (contained in) the pollard of the other (case c).

straightening map identifies v_1 and v_2 in the appropriate simplex so that the two maps coincide after all.

Finally, in the case that both codimension one degenerations are pruning degenerations, then there is no confusion about leaf sets, which are canonically identified with one another, but we must ensure that the various straightening maps agree. There are three subcases, but all are dealt with by repeated applications of Lemma A.12. See Figure 17. \square

Corollary 6.10. *For every cell K , the natural map $\mathcal{N}(K; K) \rightarrow \mathcal{N}(\text{Ch}_K(K))$ is an inclusion of pairs; restricted to the interior $\overset{\circ}{K}$ of K , the map $\mathcal{N}(\overset{\circ}{K}; K) \rightarrow \mathcal{N}(\overset{\circ}{K})$ is an isomorphism.*

Notice that Lemma 6.9 allows us to make the following definition.

Definition 6.11. Let K' be a degeneration of the cell K of arbitrary codimension. Then the map $\nabla : \mathcal{N}(\widehat{K'}; K) \rightarrow \mathcal{N}(K'; K')$ is defined to be the composition $\nabla^{\widehat{K'}} \circ \nabla^{\widehat{K_n}} \circ \dots \circ \nabla^{\widehat{K_1}}$ where $K' \hookrightarrow K_n \hookrightarrow K_{n-1} \hookrightarrow \dots \hookrightarrow K_1 \rightarrow K$ is a sequence of codimension one inclusions of cells.

By Lemma 6.9, the definition of ∇ is independent of the choice of this sequence.

6.2. Patching cohomology classes. In this section we describe how to assemble coherent collections of relative cohomology classes on K -products to give global relative cohomology classes on $\mathcal{N}(\mathcal{SD})$. We begin with a few lemmas about K -products and we return to the space $\mathcal{N}(\mathcal{SD})$ in Theorem 6.18.

Lemma 6.12. *The space N_K is homeomorphic to the total space of a $|\chi|d$ -disk bundle over the space $M^\mathcal{T}$; the space $N_K - n_K$ is homotopy equivalent to the total space of the $(|\chi|d - 1)$ -sphere bundle of the disk bundle.*

Proof. The inclusion of leaves in a tree defines a map from \mathcal{L} to \mathcal{T} , which induces a diagonal embedding of $M^\mathcal{T}$ into $M^\mathcal{L}$. Since we are working inside the injectivity

radius of M , we can project from N_K to $M^\mathcal{T}$, for example, by projecting to a single leaf coordinate in each tree. Again, because we are inside the injectivity radius, this gives N_K the structure of a disk bundle. A similar argument shows that n_K is also a disk bundle whose closure is contained in the interior of N_K . The complement $N_K - n_K$ is then homotopy equivalent to the sphere bundle of the disk bundle.

The calculation of the codimension of the embedding $M^\mathcal{T} \hookrightarrow M^\mathcal{L}$ requires an intermediate Euler characteristic argument. Recall that an oriented string diagram Γ can be written as the union of the Q_i , the T_j , and the L_i , modulo identification of vertices. The identifications occur at all leaves of T_j and at one leaf of each L_i . The overall Euler characteristic is the total number of T_j , plus the total number of L_i , minus the number of such identifications. Thus the Euler characteristic of Γ is equal to the cardinality of \mathcal{T} , minus the cardinality of \mathcal{L} .

The diagonal embedding $M^\mathcal{T} \hookrightarrow M^\mathcal{L}$ is a product over all trees T_j in \mathcal{T} of diagonal embeddings $M \hookrightarrow M^{L(T_j)}$. For each T_j the diagonal embedding has codimension $(L(T_j) - 1)d$. Therefore, the codimension of the diagonal embedding $M^\mathcal{T} \hookrightarrow M^\mathcal{L}$ is equal to $\sum_{T_j} (L(T_j) - 1)d$. By the Euler characteristic calculation above, the codimension of the diagonal embedding has codimension $|\chi|d$.

□

Remark. The space N_K is homeomorphic to the disk bundle of the normal bundle of the diagonal embedding. The space $N_K - n_K$ is homotopy equivalent to the sphere bundle of this disk bundle. This conceptually explains the connection between our construction and others [CJ02, CG04, God07a, HW14, Kup13]. Our only use of these facts in this paper will be to connect our work to theirs. See Section 8.

Lemma 6.13. *Let K be a cell of $\mathcal{SD}(\chi, k, \ell)$. Then the restriction map*

$$H^*(N_K) \rightarrow H^*(N_K - n_K)$$

is an isomorphism for $ < |\chi|d - 1$ and injective for $* = |\chi|d - 1$.*

Proof. The statement is vacuous for $|\chi|d < 1$; for $|\chi|d = 1$ it is a straightforward statement about the normal bundle of the diagonal map from the circle to the torus which is left to the reader's imagination. So for the purpose of this proof assume $|\chi|d > 1$.

The map at the level of cohomology between these two bundles is induced by a map of the respective Leray–Serre spectral sequences.

The entries of the E_2 page of the spectral sequence for N_K are $H^p(M^\mathcal{T}, H^q(\mathbb{R}^{|\chi|d}))$. The only nonzero groups occur when $q = 0$ so the sequence collapses at E_2 . On the other hand, the entries of the E_2 page of the spectral sequence for $N_K - n_K$ are $H^p(M^\mathcal{T}, H^q(S^{|\chi|d-1}))$. Therefore the nonzero entries on the E_2 page must have $q = 0$ or $q = |\chi|d - 1$. The map of spectral sequences on the E_2 page is an isomorphism onto the bottom row.

While the spectral sequence for the sphere bundle need not collapse at page E_2 , most differentials beginning with ∂_2 are zero for degree reasons. The only possibly nonzero differential is the component of $\partial_{|\chi|d}$ which goes from the entry in position $(0, |\chi|d - 1)$ to the entry in position $(|\chi|d, 0)$ on the $E_{|\chi|d}$ page.

If M is not simply connected, the cohomology groups giving the entries of the E_2 pages of both spectral sequences should be interpreted as having coefficients in local systems, but this does not affect the validity of the argument. \square

Proof. We consider the long exact sequence of the pair $(N_K, N_K - n_K)$:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & H^{m-1}(N_K) & \xrightarrow{\psi} & H^{m-1}(N_K - n_K) & & \\ & & & & \downarrow & & \\ & & & & H^m(N_K, N_K - n_K) & \xrightarrow{\varphi} & H^m(N_K) \xrightarrow{\psi} H^m(N_K - n_K) \longrightarrow \end{array}$$

Corollary 6.15. *For any cell K , the cohomology group $H^{|\mathbf{x}|d-1}(\mathcal{N}(\partial K; K))$ vanishes.*

Corollary 6.16. *Let $\widehat{K'}$ be a face of the cell K of SD . The restriction map $H^{|\mathbf{x}|d}(\mathcal{N}(\partial K; K)) \rightarrow H^{|\mathbf{x}|d}(\mathcal{N}(\widehat{K'}; K))$ is injective.*

$$H^0(\partial K) \otimes H^{|\mathbf{x}|d}(N_K, N_K - n_K) \xrightarrow{\iota^* \otimes \text{Id}} H^0(K') \otimes H^{|\mathbf{x}|d}(N_K, N_K - n_K)$$

Definition 6.17. Let $\{\tau_K\}$ be a collection of classes in $H^{|\mathbf{x}|d}(\mathcal{N}(K; K))$, as K varies over all cells of \mathcal{SD} . Assume that for each K and each codimension one degeneration K' of K that τ_K and $\tau_{K'}$ pull back to the same class in $H^{|\mathbf{x}|d}(\mathcal{N}(\widehat{K}'; K))$ under the maps induced by inclusion and ∇ respectively:

$$\iota^*(\tau_K) = \nabla^*(\tau_{K'}).$$

Then we call the collection of classes $\{\tau_K\}$ *coherent*.

Theorem 6.18. *Let $\{\tau_K\}$ be a coherent collection of cohomology classes. There exists a unique cohomology class τ in $H^{|\chi|d}(\mathcal{N}(\mathcal{SD}))$ that pulls back to τ_K under the map $\mathcal{N}(K; K) \rightarrow \mathcal{N}(\mathcal{SD})$.*

Proof. Let U be the interior of a cell K of \mathcal{SD} . Let V be the union in \mathcal{SD} of

- a subcomplex V_0 of \mathcal{SD} containing the image of the boundary of K under Ch_K but not the image of K itself, and
- the image under Ch_K of the complement of a point in the interior of K .

Then a homotopy equivalence from U to K lifts to a homotopy equivalence from $\mathcal{N}(U; K) = \mathcal{N}(U)$ to $\mathcal{N}(K; K)$. Similarly, a deformation retraction from V to V_0 lifts to a deformation retraction from $\mathcal{N}(V)$ to $\mathcal{N}(V_0)$. Note that $U \cap V$ is homotopy equivalent to ∂K . This homotopy equivalence lifts to a homotopy equivalence between $\mathcal{N}(U \cap V; K) = \mathcal{N}(U \cap V) = \mathcal{N}(U) \cap \mathcal{N}(V)$ and $\mathcal{N}(\partial K; K)$. Also note that $\mathcal{N}(U \cup V) = \mathcal{N}(U) \cup \mathcal{N}(V)$.

Assume that τ_{V_0} is an element of $H^{|\chi|d}(\mathcal{N}(V_0))$ which pulls back to $\tau_{K'}$ under the map $\mathcal{N}(K'; K') \rightarrow \mathcal{N}(V_0) \subset \mathcal{N}(\mathcal{SD})$, for each cell K' of V_0 . Since $\mathcal{N}(V)$ deformation retracts to $\mathcal{N}(V_0)$, the pullback of τ_{V_0} gives an element τ_V in $H^{|\chi|d}(\mathcal{N}(V))$ which restricts to τ_{V_0} on $\mathcal{N}(V_0)$. The restriction of τ_K to $\mathcal{N}(U)$ gives a class τ_U in $H^{|\chi|d}(\mathcal{N}(U))$. We use the Mayer–Vietoris sequence to show that τ_U and τ_V are, respectively, restrictions of a unique element $\tau_{U \cup V}$ in $H^{|\chi|d}(\mathcal{N}(U \cup V))$ to $\mathcal{N}(U)$ and $\mathcal{N}(V)$.

The relevant portion of the Mayer–Vietoris sequence is:

$$\begin{array}{ccccccc}
 & & \dots & & \xrightarrow{\psi} & H^{|\chi|d-1}(\mathcal{N}(U \cap V)) & \\
 & & \delta & & & & \\
 \hookrightarrow & H^{|\chi|d}(\mathcal{N}(U \cup V)) & \xrightarrow{\varphi} & \begin{array}{c} H^{|\chi|d}(\mathcal{N}(U)) \\ \oplus \\ H^{|\chi|d}(\mathcal{N}(V)) \end{array} & \xrightarrow{\psi} & H^{|\chi|d}(\mathcal{N}(U \cap V)) & \longrightarrow
 \end{array}$$

First we show the existence of the class $\tau_{U \cup V}$ by showing that $\psi(\tau_U, \tau_V) = 0$, that is, that τ_U and τ_V restrict to the same class in $H^{|\chi|d}(\mathcal{N}(U \cap V))$. By Corollary 6.16 it is enough to show that τ_U and τ_V restrict to the same class in $H^{|\chi|d}(\mathcal{N}(\widehat{K'}; K))$ for a face $\widehat{K'}$ in the boundary of K .

By Definition 6.8, the diagram

$$\begin{array}{ccc}
 \mathcal{N}(\widehat{K'}; K) & \xrightarrow{\quad} & \mathcal{N}(V) \\
 & \searrow \nabla^{\widehat{K'}} & \nearrow \\
 & \mathcal{N}(K'; K') &
 \end{array}$$

commutes, so the pullback of τ_V to $H^{|\chi|d}(\mathcal{N}(\widehat{K'}; K))$ is the pullback of $\tau_{K'}$ in $H^{|\chi|d}(\mathcal{N}(K'; K'))$ under the map induced by $\nabla^{\widehat{K'}}$. On the other hand, the pullback

of τ_U to $H^{|\mathbf{x}|d}(\mathcal{N}(\widehat{K'}; K))$ is given by the restriction of τ_K to $\mathcal{N}(\widehat{K'}; K)$, by the definition of τ_U . But these are equal in $H^{|\mathbf{x}|d}(\mathcal{N}(\widehat{K'}; K))$ by coherence.

Therefore $\psi(\tau_U, \tau_V) = 0$ and there exists a class $\tau_{U \cup V}$ in $H^{|\mathbf{x}|d}(\mathcal{N}(U \cup V))$ such that $\varphi(\tau_{U \cup V}) = (\tau_U, \tau_V)$.

Now we show uniqueness of $\tau_{U \cup V}$. Any two lifts $\tau_{U \cup V}$ in $H^{|\mathbf{x}|d}(\mathcal{N}(U \cup V))$ of (τ_U, τ_V) differ by an element in the image of the connecting homomorphism δ . The group $H^{|\mathbf{x}|d-1}(\mathcal{N}(U \cap V))$ vanishes by Corollary 6.15. Therefore, $\delta = 0$ and $\tau_{U \cup V}$ is unique.

Choose a total ordering of all the cells K_1, K_2, \dots, K_N of \mathcal{SD} such that the dimension of K_i is less than or equal to the dimension of K_{i+1} . We will refer to the union $\bigcup_1^i \text{Ch}_{K_i}(K_i)$ as \mathcal{SD}_i .

Now assume we have defined a class τ_i on $\mathcal{N}(\mathcal{SD}_i)$ which pulls back to τ_{K_j} on $\mathcal{N}(K_j; K_j)$ for $j \leq i$. Let p be the image under the characteristic map $\text{Ch}_{K_{i+1}}$ of an interior point of K_{i+1} . Then the space $\mathcal{N}(\mathcal{SD}_i)$ is homotopy equivalent to $\mathcal{N}(\mathcal{SD}_{i+1} - p)$.

By the Mayer–Vietoris argument above, we can then extend τ_i to τ_{i+1} . Define the class $\tau \in H^{|\mathbf{x}|d}\mathcal{N}(\mathcal{SD})$ as τ_N .

This shows the existence part of Theorem 6.18.

Now for uniqueness, let τ' be any class in $H^{|\mathbf{x}|d}(\mathcal{N}(\mathcal{SD}))$ that pulls back to τ_K for each cell K . In particular, both τ' and τ restrict to the same class on \mathcal{SD}_1 . Now assume that $\tau' \neq \tau$. Then there exists a first \mathcal{SD}_i in the sequence of subcomplexes defining τ for which the restriction of τ' is not equal to the restriction of τ . However, the class τ_i above was unique. Therefore, τ and τ' must be equal on \mathcal{SD}_i and hence on any subcomplex in the sequence. In particular τ and τ' must be equal on the entire space $\mathcal{N}(\mathcal{SD})$.

□

7. THE DIFFUSE INTERSECTION CLASS

In this section, we finally use the orientation of the d -dimensional manifold M and orientations of string diagrams. Our goal is to define a relative cohomology class in $H^{|\mathbf{x}|d}(S, S - s)$ called the *diffuse intersection class*. In this section, we define a coherent collection of Thom classes for each cell K of \mathcal{SD} , assemble them to build a global Thom class on $\mathcal{N}(\mathcal{SD})$ using Theorem 6.18, and build an evaluation map of pairs from $(S, s - s)$ to $\mathcal{N}(\mathcal{SD})$. The diffuse intersection class will be the pullback of the global Thom class under this evaluation map.

7.1. The global Thom class. Recall from the proof of Lemma 6.12 that the canonical map of leaves into trees $\mathcal{L} \rightarrow \mathcal{T}$ induces a diagonal embedding $M^{\mathcal{T}} \hookrightarrow M^{\mathcal{L}}$ whose image lies in the subspace N_K , which can be given the structure of a tubular neighborhood of $M^{\mathcal{T}}$ in $M^{\mathcal{L}}$.

Fix a string diagram Γ such that each T_j is a segment and fix an order o on Γ . There is a canonical identification of \mathcal{T} with the ordered set $\{1, \dots, |\chi|\}$. Precomposing the map induced on homology by the canonical identification of $M^{|\chi|}$ with $M^{\mathcal{T}}$ with the Eilenberg–Zilber map gives a map $H_d(M)^{\otimes |\chi|} \rightarrow H_{|\chi|d}(M^{|\chi|}) \rightarrow H_{|\chi|d}(M^{\mathcal{T}})$. We define $[M^{\mathcal{T}}]$ to be the image under this composition of the $|\chi|$ -fold tensor power of the fundamental class $[M]$ of M determined by the R -orientation of M . We define $[M^{\mathcal{L}}]$ to be the analogous element of $H_{2|\chi|d}(M^{\mathcal{L}})$ and $[N_K]$ to be its image under the composition

$$H_{2|\chi|d}(M^{\mathcal{L}}) \rightarrow H_{2|\chi|d}(M^{\mathcal{L}}, M^{\mathcal{L}} - n_K) \xrightarrow{\sim} H_{2|\chi|d}(N_K, N_K - n_K).$$

The cap product with $[N_K]$ is a Poincaré–Lefschetz duality isomorphism between $H^*(N_K, N_K - n_K)$ and $H_{2|\chi|d-*}(N_K)$. The Poincaré–Lefschetz dual to the image of $[M^{\mathcal{T}}]$ under the map induced by the diagonal embedding $M^{\mathcal{T}} \rightarrow N_K$ is an element of $H^{|\chi|d}(N_K, N_K - n_K)$. In particular, this dual element corresponds to the Thom class of the normal bundle of the diagonal embedding $M^{\mathcal{T}} \hookrightarrow M^{\mathcal{L}}$; see the remark after Lemma 6.12. If o and o' induce the same orientation on Γ , then the two dual elements in $H^{|\chi|d}(N_K, N_K - n_K)$ are equal.

Remark. In the case that \mathcal{SD} is a trivial double cover of \mathcal{SD}^u , then we could work with \mathcal{SD}^u instead of \mathcal{SD} . Additionally, if M is even-dimensional or if R has characteristic 2, then we could work with \mathcal{SD}^u instead of \mathcal{SD} .

Definition 7.1. For K a cell of \mathcal{SD} , we will define the K -Thom class ω_K , a cohomology class in $H^{|\chi|d}(\mathcal{N}(K; K))$, as follows.

If K is a 0-cell, then we denote it by p . The 0-cell p represents an oriented string diagram Γ where each T_j is a segment. Additionally, $\mathcal{N}(p; p)$ is canonically isomorphic to the pair $(N_p, N_p - n_p)$. In this case, the p -Thom class ω_p is the Poincaré–Lefschetz dual element in $H^{|\chi|d}(N_p, N_p - n_p)$ described above, which is independent of the choice of order o in the orientation class of Γ .

For K a higher dimensional cell, let \hat{p} be a 0-cell of K corresponding to the 0-cell p of \mathcal{SD} . There are maps

$$\mathcal{N}(K; K) \rightarrow \mathcal{N}(\hat{p}; K) \rightarrow \mathcal{N}(p; p).$$

The former map is induced by the projection of K onto \hat{p} ; the latter map is $\nabla^{\hat{p}}$. In this case, we define the K -Thom class at \hat{p} as the pullback of the p -Thom class ω_p along this composition. As it will turn out that this is independent of the choice of 0-cell \hat{p} , we will refer to it simply as the K -Thom class ω_K .

A priori, however, the K -Thom class ω_K depends on the choice of the vertex \hat{p} .

Lemma 7.2. *The K -Thom class at \hat{p} is independent of the choice of the 0-cell \hat{p} .*

Proof. Any two 0-cells in K can be connected by a chain of one-cells in K . Therefore, it is enough to prove that the K -Thom class at \hat{p} is equal to the K -Thom class at \hat{q} , where \hat{p} and \hat{q} are vertices of K that are connected by a one-cell $\hat{p}\hat{q}$.

The 1-cell pq of \mathcal{SD} induces an identification of the set of trees of any string diagram in the cell pq with the set of trees of any string diagram in the cell p , and likewise for

q . These identifications also respect the sets of leaves of all trees. This identification, in turn, induces an identification of N_p with N_q and of n_p with n_q . Let

$$\theta_{pq} : \mathcal{N}(p; p) \cong (N_p, N_p - n_p) \rightarrow \mathcal{N}(q; q) = (N_q, N_q - n_q)$$

be induced by these identifications. Let π_p be the composition as defined above $\mathcal{N}(K; K) \rightarrow \mathcal{N}(p; p)$ and let $\pi_q : \mathcal{N}(K; K) \rightarrow \mathcal{N}(q; q)$ be the analogous composition.

To prove that $\pi_p^*(\omega_p)$ and $\pi_q^*(\omega_q)$ are equal, we first show that the map $\theta_{pq} \circ \pi_p$ is homotopic to the map π_q and then that $\theta_{pq}^*(\omega_q) = \omega_p$.

First, to construct a homotopy, parametrize the 1-cell \widehat{pq} of K by a fixed homeomorphism $[0, 1] \rightarrow \widehat{pq}$ and consider the following map:

$$H : [0, 1] \times \mathcal{N}(K; K) \rightarrow \mathcal{N}(\widehat{pq}; K) \rightarrow \mathcal{N}(pq; pq) \rightarrow (N_{pq}, N_{pq} - n_{pq}) \rightarrow \mathcal{N}(q; q),$$

where

- the first map is induced by the parametrization of \widehat{pq} and projection from $\mathcal{N}(K; K)$ to $(N_K, N_K - n_K)$,
- the second map is $\nabla^{\widehat{pq}}$,
- the third map is projection on the second factor in both coordinates (recall that $\mathcal{N}(pq; pq)$ is $(pq \times N_{pq}, pq \times (N_{pq} - n_{pq}))$), and
- the final map uses the canonical bijection between the set of trees of pq and the set of trees of q .

It is clear that H is a homotopy between $\theta_{pq} \circ \pi_p$ and π_q .

Now, because the map $\theta_{pq} \circ \pi_p$ is homotopic to the map π_q , it suffices to show that $\theta_{pq}^*(\omega_q) = \omega_p$. Since both p and q are degenerations of the one-cell pq , there is a canonical orientation-preserving bijection between their ordered sets of trees and leaves. Then by the definition of ω_p and ω_q for 0-cells of \mathcal{SD} , they agree. \square

Lemma 7.3. *The collection $\{\omega_K\}$ of K -Thom classes is coherent.*

Proof. By Definition 7.1, the result is implied by the commutativity of the following diagram of pairs of spaces.

$$\begin{array}{ccccc}
 & & \mathcal{N}(K; K) & & \\
 & \nearrow \iota & & \searrow \pi & \\
 \mathcal{N}(\widehat{K'}; K) & \xrightarrow{\pi} & \mathcal{N}(\widehat{p}; K) & & \\
 \nabla^{\widehat{K'}} \downarrow & & \nabla^{\widehat{K'}} \downarrow & \searrow \nabla^{\widehat{p}} & \\
 \mathcal{N}(K'; K') & \xrightarrow{\pi} & \mathcal{N}(\widehat{p}; K') & \xrightarrow[\nabla^{\widehat{p}}]{} & \mathcal{N}(p; p).
 \end{array}$$

In the diagram, ι is induced by the inclusion of $\widehat{K'}$ into K and π is induced by projection onto p . In particular, both are the identity on the second factor.

The upper triangle commutes because projection to \widehat{p} is insensitive to the domain space. The rectangle commutes because the right vertical map is just the restriction of the left vertical map. The bottom right triangle commutes using Lemma 6.9 by the argument following Definition 6.11.

□

Definition 7.4. The *global Thom class* ω is the unique class in $H^{|\chi|^d}(\mathcal{N}(\mathcal{SD}))$ pulling back to ω_K in $H^{|\chi|^d}(\mathcal{N}(K; K))$ guaranteed by Theorem 6.18.

7.2. The evaluation map. Recall from Definition 4.5 that the *diffuse intersection locus* is the subspace S of $\mathcal{SD} \times LM^k$ consisting of pairs (Γ, γ) where γ is $\frac{r}{|\chi|}$ -Lipschitz with respect to Γ and the subspace s of S consists of pairs (Γ, γ) where γ is $\frac{r}{2|\chi|}$ -Lipschitz with respect to Γ .

Definition 7.5. The subspaces S_K and s_k of $K \times LM^k$ are the preimages of the subspaces S and s of $\mathcal{SD} \times LM^k$ under $\text{Ch}_K \times \text{Id}$.

Let Γ be an oriented string diagram in the cell K of \mathcal{SD} . Recall that the set of leaves \mathcal{L} has a canonical map ι into Γ . Given a map $\Theta : |\Gamma| \rightarrow M$, the pullback $\iota^*(\Theta)$ is the restriction of Θ to the image of ι in $|\Gamma|$. Recall also that $\mathcal{SD}(M)$ consists of pairs $(\Gamma, \Theta : |\Gamma| \rightarrow M)$ and the projection $\mathcal{SD}(M) \rightarrow \mathcal{SD}$ is the forgetful map $(\Gamma, \Theta) \mapsto \Gamma$. Let $K(M)$ denote the preimage in $\mathcal{SD}(M)$ of $\text{Ch}_K(K)$ under this projection map. Notice that the restriction of the map \heartsuit to S_K has image contained in $K(M)$.

Definition 7.6. The *naive K -evaluation map* is a map ev_K from S_K to $K \times M^\mathcal{L}$. It is given by:

$$S_K \xrightarrow{\heartsuit} K(M) \xrightarrow{(\text{id}, \iota^*)} K \times M^\mathcal{L}.$$

That is, on (Γ, γ) it fixes Γ and then uses \heartsuit to decide where to send \mathcal{L} .

Lemma 7.7. *The map ev_K takes S_K into $K \times N_K$.*

Proof. Let (Γ, γ) be in S_K . The length conditions on trees in Γ imply that for each component C of $\widehat{\Gamma}$, $\gamma \circ \iota$ sends the leaves of C into an r -ball in M . The straightening of C takes C to $\Delta_{L(C)}$ and the Riemannian centers of mass map $\Upsilon(\gamma \circ \iota, \cdot)$ takes $\Delta_{L(C)}$ into the r -ball containing $\gamma(\iota(L(C)))$. So for $\heartsuit(\Gamma, \gamma) = (\Gamma, \Theta)$, $\Theta : |\Gamma| \rightarrow M$ sends C into the r -ball containing $\gamma(\iota(L(C)))$. In particular, for any tree T in C , Θ sends $L(T)$ into this r -ball in M . Hence, $ev_K(\Gamma, \gamma)$ lies in $K \times N_K$. □

In order to proceed further, we use the following elementary consequence of convexity, already implicitly used in Lemma 6.6 and Figure 16.

Lemma 7.8. *Let U be a convex subset of a Riemannian manifold and let $\{P_j\}$ be a finite set of finite sets of points in U such that*

- (1) *each set of points P_j lies in an ε_j -ball in U ,*
- (2) *for any pair (i, k) there is a sequence $P_i = P_{j_1}, P_{j_2}, \dots, P_{j_n} = P_k$ such that the finite point sets P_{j_r} and $P_{j_{r+1}}$ have at least one point in common.*

Then the union of all P_j lies in a ball of radius $\sum \varepsilon_i$.

Proof. It suffices to prove the statement for $|\{P_j\}| = 2$. Let x be in $P_1 \cap P_2$ and let c_j be the center of an ε_j -ball containing P_j . Let ℓ_j be the distance from x to c_j . By convexity, there is a segment from c_1 to c_2 with length at most $\ell_1 + \ell_2 < \varepsilon_1 + \varepsilon_2$.

Pick a point y on the segment with $d(c_1, y) < \varepsilon_2$ and $d(c_2, y) < \varepsilon_1$. Then $P_1 \cup P_2$ is all within $\varepsilon_1 + \varepsilon_2$ of y . \square

Lemma 7.9. *The map ev_K takes $S_K - s_K$ into $K \times (N_K - n_K)$.*

Proof. We show that the intersection of $K \times n_K$ with the image of ev_K is contained in $ev_K(s_K)$.

Let $ev_K(\Gamma, \gamma)$ lie in $K \times n_K$ and write $\heartsuit(\Gamma, \gamma) = (\Gamma, \Theta)$. Fix a component C of the intersection graph $\widehat{\Gamma}$ made up of the trees $\{T_j\}$. A vertex of $\widehat{\Gamma}$ is in \tilde{P}_j if

- (1) it is the image of a vertex of T_j , and
- (2) it is the image of a leaf of $T_{j'}$ for any j' (possibly including j).

Let P_j be $\Theta(\tilde{P}_j) \subset M$. Since $ev_K(\Gamma, \gamma)$ is in n_K and \heartsuit relies on the straightening map and the Riemannian center of mass map, the collection P_j lies in a $|T_j| \frac{r}{4|\chi|^2}$ -ball in M . Furthermore, since the component C is connected, the second condition of Lemma 7.8 holds and we can conclude that the union $\bigcup P_j$ lies in a $\sum |T_j| \frac{r}{4|\chi|^2}$ -ball. This union contains the image of all leaves of C and the radius is at most $\frac{r}{4|\chi|}$. Thus the distance between the images of any two leaves of C is at most $\frac{r}{2|\chi|}$. This shows that (Γ, γ) is in s_K . \square

Corollary 7.10. *The naive evaluation map ev_K induces a map of pairs*

$$(S_K, S_K - s_K) \rightarrow \mathcal{N}(K; K)$$

Lemma 7.11. *Let K' be a codimension one degeneration of a cell K of \mathcal{SD} . The following diagram commutes:*

$$\begin{array}{ccc} S_K \cap (\widehat{K'} \times LM^k) & \xrightarrow{ev_K} & \widehat{K'} \times N_K \\ \downarrow & & \downarrow \nabla \\ S_{K'} & \xrightarrow{ev_{K'}} & K' \times N_{K'} \end{array}$$

where the existence of the left vertical map is guaranteed by Lemma 4.2.

Proof. If K' is a contraction degeneration of K , the statement is trivially true, so we consider only when K' is a pruning degeneration of K .

We break the diagram up into the following diagram:

$$\begin{array}{ccc} S_K \cap (\widehat{K'} \times LM^k) & \xrightarrow{ev_K} & \widehat{K'} \times N_K \\ \downarrow & \searrow \heartsuit & \downarrow \nabla \\ & K'(M) & \\ \uparrow \heartsuit & \swarrow (id, t^*) & \\ S_{K'} & \xrightarrow{ev_{K'}} & K' \times N_{K'} \end{array}$$

The bottom triangle commutes by definition of $ev_{K'}$. By Proposition 4.8, the left triangle commutes. The upper right triangle commutes because ev_K , \heartsuit and ∇ are all defined by straightening composed with the Riemannian centers of mass map.

□

We will use a relative version of this result to define a global evaluation map.

Corollary 7.12. *Let K' be a codimension one degeneration of a cell K of \mathcal{SD} . The following diagram commutes:*

$$\begin{array}{ccc} (S_K, S_K - s_K) \cap (\widehat{K'} \times LM^k, \widehat{K'} \times LM^k) & \xrightarrow{ev_K} & \mathcal{N}(\widehat{K'}; K) \\ \downarrow & & \downarrow \nabla \\ (S_{K'}, S_{K'} - s_{K'}) & \xrightarrow{ev_{K'}} & \mathcal{N}(K'; K') \end{array}$$

This finally allows us to make the following definition

Definition 7.13. The *evaluation map* $ev : (S, S - s) \rightarrow \mathcal{N}(\mathcal{SD})$ is defined on (Γ, γ) as $ev_K(\Gamma, \gamma)$, where Γ is in the cell K .

Definition 7.14. The *diffuse intersection class* Ω is the pullback to $(S, S - s)$ of the global Thom class ω on $\mathcal{N}(\mathcal{SD})$ under the evaluation map:

$$\Omega = ev^*(\omega) \in H^{|\chi|d}(S, S - s).$$

8. CONNECTIONS TO PREVIOUS WORK

Recall that the string topology construction for a cocycle W in $C^{|\chi|d}(S, S - s)$ representing the diffuse intersection class Ω is a chain map \mathcal{ST}_W from the tensor product $C_*(\mathcal{SD}) \otimes C_*(LM)^{\otimes k}$ to $C_{*-|\chi|d}(LM)^{\otimes \ell}$. By fixing a cycle α in $C_n(\mathcal{SD})$, we obtain a chain map

$$\mu_{(\alpha, W)} : C_*(LM)^{\otimes k} \rightarrow C_{*+n-|\chi|d}(LM)^{\otimes \ell}$$

which induces a map on homology

$$\bar{\mu}_{(\alpha, W)} : H_*(LM)^{\otimes k} \rightarrow H_{*+n-|\chi|d}(LM)^{\otimes \ell}$$

which depends only on the homology class of α in \mathcal{SD} .

In this section, we review the work of Cohen–Godin [CG04], who constructed a family of string topology operations

$$\mu_c : H_*(LM)^{\otimes k} \rightarrow H_{*-|\chi|d}(LM)^{\otimes \ell}$$

parameterized by a space of decorated graphs of a certain type called *marked metric chord diagrams* closely related to our string diagrams. These marked metric chord diagrams are a version of what Cohen and Godin call *Sullivan chord diagrams* with some extra data attached to them.

We isolate when marked metric chord diagrams are in fact string diagrams. Then, by choosing an arbitrary orientation, we can treat such a marked metric chord

diagram as a 0-cycle Γ in \mathcal{SD} . Then we prove Proposition 8.1, which says that our induced map on homology $\bar{\mu}_{(\Gamma, W)}$ coincides with μ_c .

Tamanoi has shown that many (but not all) of these operations are trivial [Tam10]. His methods also imply that *some* operations induced by higher degree cycles in \mathcal{SD} are trivial. One interesting question for further research is to determine precisely which higher homology classes in \mathcal{SD} induce trivial string topology operations on the homology of the loop space.

At the end of the section we also show that we recover the Batalin–Vilkovisky operator described by Chas–Sullivan in [CS99].

8.1. Homology operations induced by $H_0(\mathcal{SD})$.

Proposition 8.1. *For a marked metric chord diagram c which is the underlying string diagram of the oriented string diagram Γ , the chain map $\mu_{(\Gamma, W)}$ induces the map μ_c on homology, up to a possible sign $(-1)^d$.*

The proposition will be proved by Lemmas 8.9, 8.10, and 8.11 below. The reason for the ambiguity in the sign is as follows. Recall that our construction uses the orientation of string diagrams to define the K -Thom classes ω_K for each cell K . If the cells K and K' of \mathcal{SD} project to the same cell in \mathcal{SD}^u , then there is a canonical identification of the spaces $\mathcal{N}(K; K)$ and $\mathcal{N}(K'; K')$ which identifies the cohomology class ω_K with $(-1)^d \omega_{K'}$. Cohen–Godin do not clearly state the ordering or orientation convention they use to define their Thom classes. If M is even dimensional, then the choice of ordering or orientation is irrelevant. Thus in this case, $\mu_{(\Gamma, W)}$ induces a corresponding Cohen–Godin operation on homology with no problem. If M is odd-dimensional, $\mu_{(\Gamma, W)}$ induces the corresponding operation up to the choice of sign. Godin uses local coefficient systems to keep track of signs [God07a]; we expect that we recover her operations without ambiguity.

Note that if the string diagram Γ were in a component of \mathcal{SD} where the two-sheeted covering $\mathcal{SD} \rightarrow \mathcal{SD}^u$ was nontrivial, then $\bar{\mu}_{(\Gamma, W)}$ and $-\bar{\mu}_{(\Gamma, W)}$ would agree on all odd dimensional manifolds. In this case, Γ would induce a zero operation on the homology of the loop space of an odd dimensional manifold with coefficients in a ring where 2 was invertible. We know of no such component.

There is another potential ambiguity because it is not clearly stated in [CG04] which convention is used for identifying output boundary cycles with the standard circle. For a given boundary cycle C , one might use either the canonical map ∂_C or its reverse $\bar{\partial}_C$. We choose to use the canonical map for input boundary cycles and its reverse for output boundary cycles as yields operations that agree with [CS99]. In the absence of other evidence, we ascribe some choice along these lines to Cohen and Godin because without it, their gluing theorem (Theorem 6) fails.

Cohen and Godin construct string topology operations

$$h_*(LM)^{\otimes k} \rightarrow h_{*-|\chi|d}(LM)^{\otimes \ell}$$

for any homology theory h_* supporting an orientation of the manifold M , with coefficients in a field [CG04]. For our purposes, we consider only singular homology H_* with coefficients in a field. Their version of the space of string diagrams is the

space of marked metric chord diagrams. We adapt their definition to our notation; their original definition is in [CG04] as Definition 1 and the subsequent discussion.

Definition 8.2. A *marked metric chord diagram* is the equivalence class of a combinatorial string diagram equipped with a pseudometric structure such that

- (1) an edge has length zero if and only if it is a marking,
- (2) the union of the \tilde{T}_j contains no cycle subgraph, and
- (3) any vertex that belongs to more than one \tilde{T}_j or L_i also belongs to some $Q_{i'}$.

Γ_1 and Γ_2 are equivalent if

- (1) by pruning both Γ_1 and Γ_2 along every internal half-edge which is in some L_i one obtains isomorphic partially marked pseudometric fatgraphs with specified subfatgraphs and choices of fundamental vertices, and
- (2) under this isomorphism, corresponding output boundary cycles of Γ_1 and Γ_2 induce the same cyclic order on the subset of half-edges in that boundary cycle not in any \tilde{T}_j .

In particular, a combinatorial string diagram satisfying the appropriate properties is alone in its equivalence class unless some L_i and some \tilde{T}_j intersect, necessarily at a vertex that is also in some $Q_{i'}$. See Figure 18.




FIGURE 18. Γ_1 , Γ_2 , and their common pruning.

Cohen and Godin regard the input circles of a marked metric chord diagram c as copies of the standard circle (of variable length). They do not distinguish between the graph and its pseudometric realization explicitly; we will do so. They also introduce a metric fatgraph $S(c)$ where all length-zero edges of c have been contracted.

Cohen and Godin's primary tool for constructing a string topology operation μ_c for a marked metric chord diagram c is the Pontryagin–Thom construction for the

finite codimension subspace $\text{Maps}(|S(c)|, M)$ of LM^k . Evaluation of an element of LM^k at the leaves \mathcal{L} of the trees \mathcal{T} of c gives a map e_c from LM^k to $M^\mathcal{L}$. The subspace $\text{Maps}(|S(c)|, M)$ fits into the pullback square:

$$\begin{array}{ccc} \text{Maps}(|S(c)|, M) & \longrightarrow & LM^k \\ \downarrow \lrcorner & & \downarrow e_c \\ M^\mathcal{T} & \longrightarrow & M^\mathcal{L} \end{array}$$

where $M^\mathcal{T} \hookrightarrow M^\mathcal{L}$ is the diagonal embedding determined by c . Let ν be the normal bundle of the diagonal embedding. Then ν can be given the structure of a tubular neighborhood of the diagonal inside $M^\mathcal{L}$ by choosing an appropriate diffeomorphism. This gives $e_c^*(\nu)$ the structure of a tubular neighborhood of $\text{Maps}(|S(c)|, M)$ in LM^k ; see [God07a] for details.

For a marked metric chord diagram c with k inputs, ℓ outputs, and Euler characteristic χ , Cohen and Godin define an operation

$$\mu_c : H_*(LM)^{\otimes k} \rightarrow H_{*-|\chi|d}(LM)^{\otimes \ell}.$$

as the following composition of maps:

- (CG1) $H_*(LM)^{\otimes k} \rightarrow H_*(LM^k)$, the Künneth isomorphism,
- (CG2) next, $H_*(LM^k) \rightarrow H_*(\text{Maps}(|S(c)|, M)^\nu)$, the Thom collapse map from LM^k to the Thom space $\text{Maps}(|S(c)|, M)^\nu$ of the bundle $e_c^*(\nu)$,
- (CG3) next, $H_*(\text{Maps}(|S(c)|, M)^\nu) \rightarrow H_{*-|\chi|d}(\text{Maps}(|S(c)|, M))$, the Thom isomorphism,
- (CG4) next, $H_{*-|\chi|d}(\text{Maps}(|S(c)|, M)) \rightarrow H_{*-|\chi|d}(LM^\ell)$, induced by the pullback to outputs, and
- (CG5) finally, $H_{*-|\chi|d}(LM^\ell) \rightarrow H_{*-|\chi|d}(LM)^{\otimes \ell}$, the Künneth isomorphism.

Cohen and Godin go on to show that the operation μ_c depends only on the type (χ, k, ℓ) of the marked metric chord diagram c . Therefore, it is enough for us to consider only marked metric chord diagrams of a particularly manageable form. The definitions of marked metric chord diagrams and string diagrams give us the following lemma.

Lemma 8.3. *Let c be a marked metric chord diagram such that:*

- (1) *each L_i is disjoint from all \tilde{T}_j ,*
- (2) *every \tilde{T}_j subgraph is a length 1 segment and the images in c of all of these segments' leaves are distinct, and*
- (3) *every Q_i has total length 1.*

Then c satisfies the correct metric properties and is alone in its equivalence class and is thus a string diagram.

For the remainder of this section, assume that c is a marked metric chord diagram satisfying the conditions of Lemma 8.3 (and thus also a string diagram). Fix an arbitrary oriented string diagram Γ with c as its underlying string diagram.

It will be convenient for us to analyze the chain maps defining $\mu_{(\Gamma, W)}$ more closely. Consider the inclusion of LM^k into $\mathcal{SD} \times LM^k$ as the fiber over the point Γ . Denote

the intersection of the fiber with S by S_Γ and the intersection of the fiber with s by s_Γ ; the inclusion of the fiber induces inclusions of S_Γ into S and of $S_\Gamma - s_\Gamma$ into $S - s$.

There is a unique cell K of \mathcal{SD} such that Γ is in K but not in the boundary of K . Then the fiber over the point Γ in the space $\mathcal{N}(\mathcal{SD})$ is the pair $(N_K, N_K - n_K)$. Because Γ is an oriented string diagram where each T_j is a segment, the restriction of the global Thom class ω on $\mathcal{N}(\mathcal{SD})$ to the fiber $(N_K, N_K - n_K)$ is equal to the Thom class of the normal bundle of the embedding $M^\mathcal{T} \hookrightarrow M^\mathcal{L}$. The evaluation map sends the fiber $(S_\Gamma, S_\Gamma - s_\Gamma)$ over Γ to the fiber $(N_K, N_K - n_K)$ over Γ so the restriction of the diffuse intersection class Ω to $(S_\Gamma, S_\Gamma - s_\Gamma)$ is the pullback of the Thom class of the normal bundle under the evaluation map.

Additionally, the fiber over the point Γ in the space $\mathcal{SD}(M)$ is the mapping space $\text{Maps}(|\Gamma|, M)$ and the map \heartsuit sends the fiber over Γ to the fiber over Γ . Therefore, the map $\mu_{(\Gamma, W)}$ is given by the following composition of maps:

- (ST1) first, $C_*(LM)^{\otimes k} \rightarrow C_*(LM^k)$, the Eilenberg–Zilber shuffle map,
- (ST2) next, $C_*(LM^k) \rightarrow C_*(LM^k, LM^k - s_\Gamma)$, the usual quotient map,
- (ST3) next, $C_*(LM^k, LM^k - s_\Gamma) \rightarrow C_*(S_\Gamma, S_\Gamma - s_\Gamma)$ a chain homotopy inverse to the map induced by inclusion $(S_\Gamma, S_\Gamma - s_\Gamma) \hookrightarrow (LM^k, LM^k - s_\Gamma)$
- (ST4) next, $C_*(S_\Gamma, S_\Gamma - s_\Gamma) \rightarrow C_{*-|\chi|d}(S_\Gamma)$, the cap product with a representative of the pulled-back Thom class,
- (ST5) next, $\heartsuit_* : C_{*-|\chi|d}(S_\Gamma) \rightarrow C_{*-|\chi|d}(\text{Maps}(|\Gamma|, M))$, induced by the restriction of the map $\heartsuit : S \rightarrow \mathcal{SD}(M)$ to S_Γ ,
- (ST6) next, $C_{*-|\chi|d}(\text{Maps}(|\Gamma|, M)) \rightarrow C_{*-|\chi|d}(LM^\ell)$, induced by the pullback to outputs, and
- (ST7) finally, $C_{*-|\chi|d}(LM^\ell) \rightarrow C_{*-|\chi|d}(LM)^{\otimes \ell}$, the Alexander–Whitney map.

Since the shuffle and Alexander–Whitney maps induce the Künneth isomorphism over a field, to prove Proposition 8.1, we need only to check that the map from $H_*(LM^k)$ to $H_{*-|\chi|d}(LM^\ell)$ induced by our composition of maps (ST2) to (ST6) agrees with Cohen–Godin’s composition of maps (CG2) to (CG4), up to a possible sign. We do this step-by-step in Lemmas 8.9, 8.10, and 8.11, which focus on Cohen–Godin’s maps (CG2), (CG3), and (CG4), respectively.

We wish to realize S_Γ and s_Γ as tubular neighborhoods of $\text{Maps}(|S(c)|, M)$ in LM^k . We begin by discussing tubular neighborhoods of $M^\mathcal{T}$ in $M^\mathcal{L}$. In particular, the following definition generalizes the construction of N_K and n_K .

Definition 8.4. Let $\epsilon > 0$. Then N_K^ϵ is the subspace of $M^\mathcal{L}$ consisting of maps which, for each tree T_j , take the leaves of T_j into an ball of radius ϵ in M .

Then $N_K = N_K^r$. Furthermore, since each T_j has length 1, $n_K = N_K^{\frac{r}{4|\chi|^2}}$.

Lemma 8.5. For any $\epsilon \leq r$, the space N_K^ϵ has the structure of a tubular neighborhood of the diagonal map $M^\mathcal{T} \rightarrow M^\mathcal{L}$. Furthermore, if $\epsilon_1 < \epsilon_2$ then the closure of $N_K^{\epsilon_1}$ is contained in the interior of $N_K^{\epsilon_2}$.

Proof. The only part that does not follow directly is that N_K^ϵ is a tubular neighborhood.

The diagonal map $M^{\mathcal{T}} \hookrightarrow M^{\mathcal{L}}$ is the product over $\{T_j\}$ of the diagonal maps $M \hookrightarrow M^{L(T_j)}$. The normal bundle ν of the diagonal $M^{\mathcal{T}} \hookrightarrow M^{\mathcal{L}}$ is a product over $\{T_j\}$ of normal bundles ν_j of diagonals $M \rightarrow M^{L(T_j)}$ and we identify each ν_j with the tangent bundle TM .

The space N_K^ϵ is a product over $\{T_j\}$ of neighborhoods $N_{K,j}^\epsilon$ of the diagonals $M \rightarrow M^{L(T_j)}$ where $N_{K,j}^\epsilon$ consists of maps φ_j from $L(T_j)$ to M with images lying in ϵ balls in M .

In particular, the map φ_j sends the two leaves of the segment T_j to points x_j and y_j in the ϵ ball centered at the midpoint z_j of the unique geodesic segment joining x_j and y_j . (Note that any two of the points x_j , y_j , and z_j determine the third.) Let \exp_{z_j} be the exponential map for the tangent space to M at the point z_j and define $\phi_j : N_{K,j}^\epsilon \rightarrow \nu_j$ as $\phi_j(x_j, y_j) = \exp_{z_j}^{-1}(x_j)$, identifying ν_j with TM , the tangent bundle of M . Then ϕ_j maps the image of the diagonal to the zero section of ν_j , and is a diffeomorphism onto its image which is the ϵ -disk bundle of ν_j .

Finally, define $\phi : N_K^\epsilon \rightarrow \nu$ as the product over $\{T_j\}$ of diffeomorphisms ϕ_j . Then ϕ maps the image of the diagonal to the zero section of ν and is a diffeomorphism onto its image which is a convex neighborhood of the zero section of ν .

□

Corollary 8.6. *Let $0 < \epsilon_1 < \epsilon_2 \leq r$; let $0 < s_1 < s_2$. Let ν be the normal bundle of the embedding $M^{\mathcal{T}} \rightarrow M^{\mathcal{L}}$; let $B_s \nu$ be the radius s disk bundle of ν .*

Then there is an isomorphism of fiber bundle pairs between $(B_{s_2} \nu, B_{s_2} \nu - B_{s_1} \nu)$ and $(N_K^{\epsilon_2}, N_K^{\epsilon_2} - N_K^{\epsilon_1})$.

Lemma 8.7. *The subspace S_Γ (respectively s_Γ) of LM^k is the pullback of $N_K^{\frac{r}{|\mathbf{x}|}}$ (respectively $N_K^{\frac{r}{2|\mathbf{x}|}}$) under the evaluation map e_c .*

Proof. This is true roughly by definition. The conditions from Lemma 8.3 imply that each tree T_j is a separate component of the intersection graph of Γ . □

Corollary 8.8. *Let $0 < s_1 < s_2$. The isomorphism of fiber bundle pairs from Corollary 8.6 pulls back to an isomorphism fiber bundle pairs between*

$$(e_c^*(B_{s_2} \nu), e_c^*(B_{s_2} \nu - B_{s_1} \nu)) = (e_c^*(B_{s_2} \nu), e_c^*(B_{s_2} \nu) - e_c^*(B_{s_1} \nu))$$

and

$$(S_\Gamma, S_\Gamma - s_\Gamma).$$

We will use the notation f for the isomorphism from Corollary 8.8.

Now to prove that our operations and those of Cohen–Godin agree, we show that each of a sequence of three diagrams commute.

Lemma 8.9. *The following diagram commutes:*

$$\begin{array}{ccc}
 H_*(LM^k) & \xrightarrow{(CG2)} & H_*(\text{Maps}(|S(c)|, M)^\nu) \\
 (ST2) \downarrow & \nearrow & \uparrow \\
 H_*(LM^k, LM^k - s_\Gamma) & & \\
 (ST3) \downarrow \sim & & \\
 H_*(S_\Gamma, S_\Gamma - s_\Gamma) & \xrightarrow{f_*} & H_*(e_c^*(B_{s_2}\nu), e_c^*(B_{s_2}\nu) - e_c^*(B_{s_1}\nu))
 \end{array}$$

where the lower right horizontal map is induced by the diffeomorphism of Corollary 8.8 and the right vertical map is the isomorphism on homology that exists because the fiber bundle pair is a cofibration.

Proof. This is true by definition of the Thom collapse map. There is a quotient to relative homology followed by excision from the upper left corner to the bottom right corner which commutes by construction with the lower left maps and by definition with the upper right maps. \square

Now note that since Γ is merely an oriented version of c , there is a map at the level of pseudometric realizations $|\Gamma| = |c| \rightarrow |S(c)|$ given by contracting images of trees. This map induces an inclusion i of $\text{Maps}(|S(c)|, M)$ into $\text{Maps}(|\Gamma|, M)$ as maps which are constant on images of trees.

Lemma 8.10. *The following diagram commutes:*

$$\begin{array}{ccccc}
 H_*(\text{Maps}(|S(c)|, M)^\nu) & \xrightarrow{(CG3)} & H_{*-|\chi|d}(\text{Maps}(|S(c)|, M)) & & \\
 \uparrow & & & & \downarrow i_* \\
 H_*(S_\Gamma, S_\Gamma - s_\Gamma) & \xrightarrow{(ST4)} & H_{*-|\chi|d}(S_\Gamma) & \xrightarrow{(ST5)} & H_{*-|\chi|d}(\text{Maps}(|\Gamma|, M))
 \end{array}$$

where the left vertical map is induced by the diffeomorphism of Corollary 8.8 and the right vertical map is induced by the inclusion i .

Proof. The Thom isomorphism $H_*(\text{Maps}(|S(c)|, M)^\nu) \xrightarrow{(CG3)} H_*(\text{Maps}(|S(c)|, M))$ splits as a composition of two isomorphisms:

- (CG3)' first, the isomorphism from $H_*(\text{Maps}(|S(c)|, M)^\nu)$ to $H_{*-|\chi|d}(e_c^*(\nu))$ given by capping with the Thom class of the bundle $e_c^*(\nu)$, and
- (CG3)'' then, the isomorphism from $H_{*-|\chi|d}(e_c^*(\nu))$ to $H_{*-|\chi|d}(\text{Maps}(|S(c)|, M))$ induced by projection p from the total space to the base space of the pulled back normal bundle.

Therefore, the diagram above may be rewritten as two adjacent squares. The square on the left commutes, possibly up to sign $(-1)^d$, again by Corollary 8.8. A priori, the Thom class on the bottom and the Thom class on the top are Thom classes of pulled back fiber bundle pairs $(N_K^{\epsilon_2}, N_K^{\epsilon_2} - N_K^{\epsilon_1})$ for two different choices of (ϵ_1, ϵ_2) , but there is a map of pairs between these two choices which makes everything

commute.

$$\begin{array}{ccccc}
 H_*(\text{Maps}(|S(c)|, M)^\nu) & \xrightarrow{(CG3)'} & H_{*-|\chi|d}(e_c^*(\nu)) & \xrightarrow{(CG3)''} & H_{*-|\chi|d}(\text{Maps}(|S(c)|, M)) \\
 \uparrow f_* & & \uparrow f_* & & \downarrow i_* \\
 H_*(S_\Gamma, S_\Gamma - s_\Gamma) & \xrightarrow{(ST4)} & H_{*-|\chi|d}(S_\Gamma) & \xrightarrow{(ST5)} & H_{*-|\chi|d}(\text{Maps}(|\Gamma|, M))
 \end{array}$$

The square on the right is induced by a diagram in spaces:

$$\begin{array}{ccc}
 e_c^*(\nu) & \xrightarrow{p} & \text{Maps}(|S(c)|, M) \\
 \uparrow f & & \downarrow i \\
 S_\Gamma & \xrightarrow{\heartsuit} & \text{Maps}(|\Gamma|, M)
 \end{array}$$

This diagram does not commute on the nose, but it does commute up to homotopy. Roughly, for a fixed element γ of S_Γ , $\heartsuit(\gamma)$ extends γ to a map from $|\Gamma|$ to M by mapping trees to short geodesic segments. Following the same logic as in the proof of Lemma 6.12, there is a homotopy from such a map to one which is constant on each tree. A consistent choice of such a homotopy for all points γ in S_Γ gives rise to a homotopy from \heartsuit to the composition $i \circ p \circ f$. Therefore, the induced diagram on the homology of these spaces commutes. \square

Lemma 8.11. *The following diagram commutes:*

$$\begin{array}{ccc}
 H_{*-|\chi|d}(\text{Maps}(|S(c)|, M)) & & \\
 \downarrow i_* & \searrow (CG4) & \\
 & & H_{*-|\chi|d}(LM^\ell) \\
 & \nearrow (ST6) & \\
 H_{*-|\chi|d}(\text{Maps}(|\Gamma|, M)) & &
 \end{array}$$


Proof. We show that the following diagram in spaces commutes up to homotopy:

$$\begin{array}{ccc}
 \text{Maps}(|S(c)|, M) & & \\
 \downarrow i & \searrow \rho_{out} & \\
 & & LM^\ell \\
 & \nearrow \rho_{out} & \\
 \text{Maps}(|\Gamma|, M) & &
 \end{array}$$

where the two maps ρ_{out} to LM^ℓ are given by pulling back maps from $|S(c)|$ and $|\Gamma|$ respectively to output circles.

Consider a single output circle of Γ , with $S^1 \rightarrow |\Gamma|$. The composition of maps $S^1 \rightarrow |\Gamma| \rightarrow |S(c)|$ contracts distinct subintervals of S^1 which are preimages of trees of Γ . The corresponding output of $S(c)$, $S^1 \rightarrow |S(c)|$ does not have such contracted subintervals. Therefore, the following diagram commutes up to a homotopy

parametrizing the contraction of the subintervals and concomitant rescaling:



Postcomposing this homotopy with a map θ in $\text{Maps}(|S(c)|, M)$ gives a homotopy of the two pullbacks to this output, that is, a homotopy of the two maps from S^1 to M . For this fixed θ , the disjoint union over all ℓ outputs of these homotopies gives a map from $\sqcup_\ell S^1 \times I \rightarrow M$ which gives rise to a homotopy of maps in the diagram. Therefore, the diagram of mapping spaces commutes up to homotopy. \square

By placing these three commutative diagrams side by side, we see that both constructions induce the same operations from $H_*(LM^k)$ to $H_{*-|\chi|d}(LM^\ell)$ and the proposition follows.

8.2. The Batalin–Vilkovisky operator. In this final section we shall show that we recover the Batalin–Vilkovisky operator defined by Chas–Sullivan [CS99].

Consider the map $\mathfrak{r} : S^1 \times LM \rightarrow LM$ given by $\mathfrak{r}(s, \gamma)(t) = \gamma(s + t \pmod{1})$. This map induces a map on homology $\Delta : H_*(LM) \rightarrow H_{*+1}(LM)$ given by

$$\Delta(\alpha) = \mathfrak{r}_*(\eta \times \alpha)$$

where η is the class of the defining map $[0, 1] \rightarrow S^1$, viewed as a 1-chain.

At the chain level, we may make the following definition:

Definition 8.12. The *chain-level Batalin–Vilkovisky operator* Δ is given by the composition

$$C_i(LM) \xrightarrow{\eta \otimes} C_1(S^1) \otimes C_i(LM) \xrightarrow{\text{EZ}} C_{i+1}(S^1 \times LM) \xrightarrow{\mathfrak{r}_*} C_{i+1}(LM).$$

Now consider our space $\mathcal{SD}(0, 1, 1)$ of string diagrams of Euler characteristic 0 with one input and one output. A string diagram Γ in that space necessarily has a single Q_1 and a single L_1 with empty $\{\hat{T}_j\}$ set. For the input *in* and the output *out* of such a diagram Γ , the maps ∂_{in} and $\bar{\partial}_{out}$ provide explicit identifications of the pseudometric realization $|\Gamma|$ of Γ with the standard circle. On such a string diagram, an ordering or orientation is no data at all. There is an explicit cellular identification p of the space $\mathcal{SD}(0, 1, 1)$ with the standard circle unit circle S^1 given by the position on Q_1 of the vertex of L_1 , relative to the vertex of Q_1 . See Figure 19. The string diagram for which these vertices coincide corresponds to the 0-cell; string diagrams for which these vertices are distinct correspond to points in the interior of the 1-cell. An explicit formula for the identification is given by $p(\Gamma) = \partial_{in}^{-1} \bar{\partial}_{out}(0)$.

Proposition 8.13. Our string topology operation $\mu_{p^{-1}\eta, W} : C_*(LM) \rightarrow C_{*+1}(LM)$ coincides with the chain-level Batalin–Vilkovisky operator Δ .




FIGURE 19. The space $\mathcal{SD}(0, 1, 1)$ and two string diagrams; one corresponding to the 0-cell and one corresponding to a point in the interior of the 1-cell.

Proof. Because there are no T_j in any string diagram in $\mathcal{SD}(0, 1, 1)$, the string topology construction is significantly simpler. The spaces S and s are both equal to $\mathcal{SD}(0, 1, 1) \times LM$, so our relative chains are absolute chains and our excision map is the identity.

The embedding of $M^{\mathcal{T}}$ in $M^{\mathcal{L}}$ is the constant map on the point and its Poincaré–Lefschetz dual is the degree 0 cohomology class represented by the constant function 1. Chasing the definitions, we see that the diffuse intersection class Ω is also the class in degree 0 cohomology of the constant function 1. In degree 0, cocycle representatives are unique, so W must be the constant function 1 and capping with it is the identity.

The map \heartsuit uses the canonical map ∂_{in} for the input boundary cycle and takes the pair (Γ, γ) to the pair $(\Gamma, \gamma \circ \partial_{in}^{-1})$ in $\mathcal{SD}(0, 1, 1)(M)$.

Then the entire composition $(\rho_{in})_!$ described in Definition 5.1 boils down to the Eilenberg–Zilber shuffle map $C_1(\mathcal{SD}(0, 1, 1)) \otimes C_i(LM) \rightarrow C_{i+1}(\mathcal{SD}(0, 1, 1) \times LM)$ followed by the map induced by \heartsuit .

Next, the map $(\rho_{out})_*$ described in Definition 5.2 is induced by restriction to the single output using ∂_{out} (the Alexander–Whitney map is the identity). Then our map is the composition that runs along the left side of the following commutative diagram, whereas the chain level Batalin–Vilkovisky operator is the composition

that runs along the right:

$$\begin{array}{ccccc}
& & C_i(LM) & & \\
& \swarrow p^{-1}\eta \otimes & & \searrow \eta \otimes & \\
C_1(\mathcal{SD}(0, 1, 1)) \otimes C_i(LM) & \xrightarrow{p_* \otimes \text{id}} & C_1(S^1) \otimes C_i(LM) & & \\
\downarrow \text{EZ} & & \downarrow \text{EZ} & & \\
C_{i+1}(\mathcal{SD}(0, 1, 1) \times LM) & \xrightarrow{(p \times \text{id})_*} & C_{i+1}(S^1 \times LM) & & \\
& \searrow (\rho_{out})_* \circ \heartsuit_* & & \swarrow \mathfrak{r}_* & \\
& & C_{i+1}(LM) & &
\end{array}$$

The upper triangle and square clearly commute. To see commutativity of the bottom triangle, observe that it is induced by a diagram of spaces

$$\begin{array}{ccc}
\mathcal{SD}(0, 1, 1) \times LM & \xrightarrow{p \times \text{id}} & S^1 \times LM \\
& \searrow \rho_{out} \circ \heartsuit & \swarrow \mathfrak{r} \\
& & LM.
\end{array}$$

The two compositions are as follows:

$$\begin{aligned}
(\rho_{out} \circ \heartsuit)(\Gamma, \gamma)(t) &= \gamma(\partial_{in}^{-1} \circ \bar{\partial}_{out}(t)) \\
\mathfrak{r}(p(\Gamma), \gamma)(t) &= \gamma(t + \partial_{in}^{-1} \bar{\partial}_{out}(0)).
\end{aligned}$$

But $\partial_{in}^{-1} \bar{\partial}_{out} : S^1 \rightarrow S^1$ is always a rotation so these coincide and the diagram commutes. \square

APPENDIX A. STRAIGHTENING OF SHORT-BRANCHED TREES

The purpose of this appendix is to prove Proposition 3.1. We will do this constructively by defining a straightening map str which satisfies the conditions of the proposition. Namely, given a short-branched tree T , the map str from the pseudometric realization $|T|$ of T to the simplex $\Delta_{L(T)}$ spanned by the leaves of T takes each leaf to itself. Further, if T has an internal edge of length zero or a prunable branch, then the straighten map behaves well with respect to contracting the edge or pruning the branch.

To assign a point in $\Delta_{L(T)}$ to a point v of $|T|$ it suffices to give barycentric coordinates. This means that for every leaf $w \in L(T)$ we give v a nonnegative coordinate $a(v, w)$ corresponding to w such that the sum over all leaves of the coordinates $a(v, w)$ is 1.

Our formula for $a(v, w)$ includes expressions involving lengths of branches of trees. Because we are working in $|T|$ rather than in T itself, we have definitions analogous to those given previously. In particular, we define versions of *vertices*, *edges*, *length*, and *leaf length* in the metric space $|T|$ rather than in T itself.

Fix a point v of $|T|$.

- (1) A *v-zero cell* of $|T|$ is a point in the image of a vertex of T or v itself, if v is not in the image of any vertex of T .
- (2) A *v-one cell* of $|T|$ is a subspace of $|T|$ homeomorphic to a closed interval whose boundary consists of two *v-zero cells* and which contains no other *v-zero cells*.
- (3) The *length* of a *v-one cell* is the distance in $|T|$ between its two boundary *v-zero cells*.
- (4) Given a subspace of $|T|$ consisting of *v-one cells*, its *length* is the sum of the lengths of the constituent *v-one cells*.
- (5) Given a pair of distinct *v-zero cells* v_i and v_j of $|T|$ the subspace $T(v_i, v_j)$ is the closure in $|T|$ of the connected component of $|T| - v_i$ containing v_j .
- (6) Given such a subspace $T(v_i, v_j)$ of $|T|$, its *leaf length* is the number of leaves of T whose image in $|T|$ lies in $T(v_i, v_j) - v_i$.
- (7) Given such a subspace $T(v_i, v_j)$ of $|T|$, its *deviation* $D(v_i, v_j)$ is the leaf length of $T(v_i, v_j)$ minus the length of $T(v_i, v_j)$.

Now we prove a sequence of lemmas giving bounds on the values of the deviations. We will use these lemmas to define barycentric coordinate $a(v, w)$.

Lemma A.1. *Given a pair distinct v-zero cells v_i and v_j of $|T|$, the deviation $D(v_i, v_j)$ of $T(v_i, v_j)$ is greater than or equal to zero.*

Proof. Consider the subtree of T consisting of edges of T whose images in $|T|$ intersect $T(v_i, v_j) - v_i$. At most one leaf of this subtree is not a leaf of T . In the case that this subtree is a branch of T or all of T itself, we name this subtree $\hat{T}(v_i, v_j)$. Otherwise, the subtree has one leaf which is a bivalent vertex of T . In this case, we add to the subtree the other edge of T adjacent to this bivalent vertex. We continue adding edges in this way until we reach a non-bivalent vertex of T and name the resulting subtree $\hat{T}(v_i, v_j)$. In the case that this last vertex is at least trivalent in T , the subtree $\hat{T}(v_i, v_j)$ is a branch of T . In the case that this last vertex is a leaf of T , the subtree $\hat{T}(v_i, v_j)$ is T itself.

The leaf length of $\hat{T}(v_i, v_j)$ is equal to the leaf length of $T(v_i, v_j)$. The length of $\hat{T}(v_i, v_j)$ is greater than or equal to the length of $T(v_i, v_j)$. Because T is a short-branched tree, the length of $\hat{T}(v_i, v_j)$ is less than or equal to its leaf length. Thus the length of $T(v_i, v_j)$ is less than or equal to its leaf length. See Figure 20.

□

Lemma A.2. *Let T be a short-branched tree. Let v_i be a v-zero cell of $|T|$ and let $\{v_j\}$ be a set of points, one from each component of $|T| - v_i$. Then*

$$\sum_j D(v_i, v_j) = \begin{cases} 0 & \text{if } v_i \text{ is the image of a leaf of } T \\ 1 & \text{otherwise.} \end{cases}$$

Figure 21 shows the space $|T|$, together with the *v-zero cell* v_i and the set $\{v_j, v'_j, v''_j\}$, each in a different component of $|T| - v_i$. The space $|T|$ is the union of subspaces $T(v_i, v_j)$, $T(v_i, v'_j)$, and $T(v_i, v''_j)$




FIGURE 20. Construction of the branch $\hat{T}(v_i, v_j)$ where $v_i = v$ is not the image of a vertex of T . The first picture shows the v -zero and v -one cells of $|T|$ and the second shows the subspace $T(v_i, v_j)$ of $|T|$. The third picture shows the initial subtree of T and the fourth shows its completion to the branch $\hat{T}(v_i, v_j)$ of T .




FIGURE 21. The space $|T|$ as the union of subspaces $T(v_i, v_j)$, $T(v_i, v'_j)$, and $T(v_i, v''_j)$.

Proof. By definition,

$$\sum_j D(v_i, v_j) = \sum_j \text{leaf length of } T(v_i, v_j) - \sum_j \text{length of } T(v_i, v_j).$$

The first sum on the right is equal to the number of leaves of T whose images lie in $|T| - v_i$. The second sum on the right is the length of $|T|$, which is the length of T . Because T is a short-branched tree, the length of T is equal to one less than the number of leaves of T . Thus, if v_i is itself the image of a leaf, the two sums on the right are equal; if v_i is not the image of a leaf, the first sum is one greater than the second. \square

Lemma A.3. *Given a pair of distinct v -zero cells v_i and v_j of $|T|$, the deviation $D(v_i, v_j)$ of $T(v_i, v_j)$ is strictly less than one.*

Proof. If v_i is the image of a leaf of T , then the previous two lemmas show that $D(v_i, v_j) = 0$.

If v_i is not the image of a leaf of T , then assume $D(v_i, v_j) = 1$. Then for any v'_j in a different component of $|T| - v_i$ than the one containing v_j , $D(v_i, v'_j)$ must be 0 by the previous two lemmas. Let e be the v -one cell of $|T|$ adjacent to v_i in $T(v_i, v_j)$ and consider the subspace of $|T|$ given by the union of the $T(v_i, v'_j)$, for all such v'_j and e . If v'_i is the other v -zero cell of e then this subspace is equal to $T(v'_i, v_i)$. The leaf length of $T(v'_i, v_i)$ is equal to the sums of the leaf lengths of the $T(v_i, v'_j)$ and the length of $T(v'_i, v_i)$ is equal to the sums of the lengths of the $T(v_i, v'_j)$, plus the length of e . Thus, $D(v'_i, v_i)$ is equal to the sum of the $D(v_i, v'_j)$, minus the length of e and $D(v'_i, v_i)$ is equal to negative the length of e . But e has positive length and $D(v'_i, v_i)$ must be nonnegative by Lemma A.1, so we arrive at a contradiction. \square

Lemma A.4. *Given a v -one cell e of $|T|$ containing the v -zero cells v_i and v_j , the sum of deviations $D(v_i, v_j) + D(v_j, v_i)$ is equal to one minus the length of e .*

Proof. The sum of the leaf lengths of $T(v_i, v_j)$ and $T(v_j, v_i)$ is equal to the total number of leaves of T . The sum of the lengths of $T(v_i, v_j)$ and $T(v_j, v_i)$ is equal to the length of $|T|$ plus the length of e . But the length of $|T|$ is equal to the length of T , which is equal to one less than the number of leaves of T . Therefore the sum of deviations $D(v_i, v_j) + D(v_j, v_i)$ is equal to one minus the length of e . \square

Corollary A.5. *Given a v -one cell e of $|T|$ containing the v -zero cells v_i and v_j , $\frac{D(v_i, v_j)}{1 - D(v_j, v_i)}$ is well-defined and is between 0 and 1.*

We are now ready to define the barycentric coordinate $a(v, w)$, which uses a sequence of v -zero cells in $|T|$. Given a point v of $|T|$ and a leaf w of T , any path in $|T|$ from v to the image $|w|$ of w passes through a finite sequence of v -zero cells. In particular, there is a unique such sequence $v = v_1, v_2, \dots, v_n = |w|$ of *distinct* v -zero cells.




FIGURE 22. A path from a point v to a leaf w in the short-branched tree T

Definition A.6. The *barycentric coordinate* $a(v, w)$ is given by the formula

$$a(v, w) = \prod_{i=1}^{n-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)}.$$




FIGURE 23. A short-branched tree T and its straightening in $\Delta_{L(T)}$

Lemma A.7. *The barycentric coordinate is continuous with respect to the first coordinate.*

Proof. Fix an oriented edge \vec{e} of T with length $\ell(\vec{e}) > 0$. In general, the factors $\frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)}$ defining $a(v, w)$ vary continuously—and in fact, all but the first one of them are constant—as v moves around in the image of $(0, \ell(\vec{e})) \times \{\vec{e}\}$ in $|T|$. Therefore it suffices to consider a sequence of points $\{v^k\}$ in the image of $(0, \ell(\vec{e})) \times \{\vec{e}\}$ in $|T|$ converging to the point v^∞ , the image of $\{0\} \times \{\vec{e}\}$ in $|T|$.

Unless w is in $T(v^k, v^\infty)$, the factors vary continuously and therefore $a(v^k, w)$ converges to $a(v^\infty, w)$.

If w is in $T(v^k, v^\infty)$, then in the limit the first factor $\frac{D(v^k, v^\infty)}{1 - D(v^\infty, v^k)}$ defining $a(v^k, w)$ disappears in the product defining $a(v^\infty, w)$. By Lemma A.4, the sum of deviations $D(v^k, v^\infty) + D(v^\infty, v^k)$ is equal to one minus the length of the v^k -one cell e joining v^k and v^∞ . When v^k approaches v^∞ , the length of e approaches zero. Therefore the first factor $\frac{D(v^k, v^\infty)}{1 - D(v^\infty, v^k)}$ approaches one and $a(v^k, w)$ converges to $a(v^\infty, w)$. \square

Lemma A.8. *For any point v in $|T|$, the sum over all leaves w in $L(T)$ of the barycentric coordinates $a(v, w)$ is 1.*

Proof. Fix a v -zero cell \hat{v} of $|T|$.

If v is the image of a leaf of T , then the first factor defining $a(v, w)$ is zero unless $v = |w|$. If $v = |w|$, the product is empty, so $a(v, w) = 1$.

Therefore, assume that v is not the image of a leaf in T . Consider the set of leaves $\{w_i\}$ such that every path from v to $|w_i|$ in $|T|$ passes through \hat{v} . Define $a(v, \hat{v})$ to

be the sum $\sum_{w \in \{w_i\}} a(v, w)$. Notice that

$$\sum_w a(v, w) = \sum_{\hat{v}} a(v, \hat{v}),$$

where the sum on the left is over all leaves w of T and the sum on the right is over all vertices \hat{v} which are adjacent to v in T , so it suffices to prove that $\sum_{\hat{v}} a(v, \hat{v}) = 1$.

In fact, we will prove a more general statement than this. We will show that for a general v -zero cell \hat{v} of $|T|$,

$$a(v, \hat{v}) = D(v_{n-1}, \hat{v}) \prod_{i=1}^{n-2} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)}$$

where $v = v_1, v_2, \dots, v_n = \hat{v}$ is the sequence of v -zero cells in $|T|$ between v to \hat{v} . This will prove what we want, that $\sum_{\hat{v}} a(v, \hat{v}) = 1$ for \hat{v} adjacent to v , since in this case $a(v, \hat{v}) = D(v, \hat{v})$. The result then follows by Lemma A.2.

We now prove the general statement, that by induction over the vertices \hat{v} .

In the base case, \hat{v} is the image of a leaf of T , so the formula holds by our original definition of $a(v, w)$ since all deviations $D(\hat{v}, v')$ are equal to zero.

Now assume that \hat{v} is not the image of a leaf of T and that this formula holds for $a(v, \tilde{v})$, where \tilde{v} is a v -zero cell sharing a v -one cell with \hat{v} and every path from v to \tilde{v} passes through \hat{v} . For such \tilde{v} , let $v = v_1, \dots, \hat{v} = v_n, \tilde{v} = v_{n+1}$ be the sequence of v -zero cells between v and \tilde{v} . So by assumption

$$a(v, \tilde{v}) = D(\hat{v}, \tilde{v}) \prod_{i=1}^{n-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)}.$$

We need to sum over all choices of \tilde{v} as above. Doing, this we get

$$a(v, \hat{v}) = \sum_{\tilde{v}} a(v, \tilde{v}) = \left(\prod_{i=1}^{n-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)} \right) \sum_{\tilde{v}} D(\hat{v}, \tilde{v}).$$

By Lemma A.2, the sum $\sum_{\tilde{v}} D(\hat{v}, \tilde{v}) = 1 - D(\hat{v}, v_{n-1})$.

□

Corollary A.9. *The barycentric coordinate is a map from $|T|$ to $\Delta_{L(T)}$ that sends each leaf to itself.*

Definition A.10. The *straightening* of the short-branched tree T with leaf set $L(T)$ is the map from $|T|$ to $\Delta_{L(T)}$ defined by barycentric coordinates.

By Corollary A.9, the straightening map str satisfies the first condition of Proposition 3.1. It remains to show that it satisfies the second and third conditions. Namely, the straightening map degenerates nicely with respect to contraction degenerations (Lemma A.11) and with respect to pruning degenerations (Lemma A.12).

Lemma A.11. *Let e be an internal edge of T or an external edge of T whose internal vertex is bivalent. Assume e has length zero. Let $|T| \rightarrow |T/e|$ be the isomorphism induced by the contraction of the edge e . Then the following diagram commutes.*

$$\begin{array}{ccc} |T| & \xrightarrow{\text{str}} & \Delta_{L(T)} \\ \cong \downarrow & & \downarrow \cong \\ |T/e| & \xrightarrow{\text{str}} & \Delta_{L(T/e)} \end{array}$$

Proof. For any point v of $|T|$, the isomorphism $|T| \rightarrow |T/e|$ preserves v -zero cells and lengths of v -one cells. Therefore the barycentric coordinates of Definition A.6 coincide in the two cases. \square

Lemma A.12. *Let T_h be a prunable branch of T with pollard T^h . Since $L(T^h)$ is a subset of $L(T)$, there is a natural inclusion of $\Delta_{L(T^h)}$ in $\Delta_{L(T)}$. Since every leaf of T_h except $s(h)$ is also a leaf of T , assigning a point in $\Delta_{L(T)}$ to $s(h)$ yields a linear inclusion of $\Delta_{L(T_h)}$ in $\Delta_{L(T)}$. Here, since $s(h)$ is also a vertex of T^h , it has an image point $|s(h)|$ in the pseudometric realization $|T^h|$. Thus we assign to $s(h)$ the image of $|s(h)|$ under the straightening map of T^h . Then the following diagram commutes.*

$$\begin{array}{ccc} |T^h| \sqcup |T_h| & \xrightarrow{\text{str}} & \Delta_{L(T^h)} \sqcup \Delta_{L(T_h)} \\ \rho \downarrow & & \downarrow \\ |T| & \xrightarrow{\text{str}} & \Delta_{L(T)} \end{array}$$

See Figure 24 for a picture of the inclusion of $\Delta_{L(T_h)}$ in $\Delta_{L(T)}$.

Proof. Let v be a point in $|T^h| \sqcup |T_h|$ and let w be a leaf of T . We prove different cases, which depend on whether the preimage of $|w|$ in $|T^h| \sqcup |T_h|$ under ρ lies in the same component as v or not.

Case 1. If the point v is in $|T^h|$ and there is a leaf w^h of T^h with $\rho(|w^h|) = |w|$, then the choice of leaf w^h is unique. Let $T^h(v_i, v_j)$ be a subspace of $|T^h|$ whose deviation appears in the expression for $a(v, w^h)$ in $\Delta_{L(T^h)}$. If $|s(h)|$ is not a point in $T^h(v_i, v_j) - v_i$, then the deviation of the isomorphic image $T(v_i, v_j) = \rho(T^h(v_i, v_j))$ appears in the expression for $a(v, w)$ in $\Delta_{L(T)}$. On the other hand, if $|s(h)|$ is a point in $T^h(v_i, v_j) - v_i$, then in the expression for $a(v, w)$ in $\Delta_{L(T)}$, the subspace $T(v_i, v_j)$ of $|T|$ corresponding to $T^h(v_i, v_j)$ is instead the image under ρ of the union of $(|T_h|)$ with $T^h(v_i, v_j)$. Because T_h is prunable, the deviation $D(v_i, v_j)$ is equal to $D^h(v_i, v_j)$. Each factor of $a(v, w^h)$ in $\Delta_{L(T^h)}$ is equal to the corresponding factor of $a(v, w)$ in $\Delta_{L(T)}$, so $a(v, w^h) = a(v, w)$.

Case 2. If the point v is in $|T^h|$ and there is no leaf w^h of T^h with $\rho(|w^h|) = |w|$, then notice by Lemma A.8, the sum over leaves of T^h of $a(v, w) = 1$, and likewise for the sum over leaves of T . Then the argument for Case 1 also shows that in this case $a(v, w) = 0$.

Case 3. If the point v is in $|T_h|$ and there is no leaf w^h of T^h with $\rho(|w^h|) = |w|$, then there is a unique leaf w_h of T_h so that $\rho(|w_h|) = |w|$. Now this case is analogous to the first case.

Let $T_h(v_i, v_j)$ be a subspace of $|T_h|$ whose deviation appears in the expression for $a(v, w_h)$ in $\Delta_{L(T_h)}$. If $|s(h)|$ is not a point in $T_h(v_i, v_j) - v_i$, then the deviation of the isomorphic image $T(v_i, v_j) = \rho(T_h(v_i, v_j))$ appears in the expression for $a(v, w)$ in $\Delta_{L(T)}$. On the other hand, if $|s(h)|$ is a point in $T_h(v_i, v_j) - v_i$, then in the expression for $a(v, w)$ in $\Delta_{L(T)}$, the subspace $T(v_i, v_j)$ of $|T|$ corresponding to $T_h(v_i, v_j)$ is instead the image under ρ of the union of $(|T^h|)$ with $T_h(v_i, v_j)$. As in the previous case, prunability of T_h implies that the deviation $D(v_i, v_j)$ is equal to $D_h(v_i, v_j)$. Each factor of $a(v, w_h)$ in $\Delta_{L(T_h)}$ is equal to the corresponding factor of $a(v, w)$ in $\Delta_{L(T)}$, so $a(v, w_h) = a(v, w)$.

Case 4. If the point v is in $|T_h|$ and there is a leaf w^h of T^h with $\rho(|w^h|) = |w|$, then any path from v to $|w|$ in $|T|$ must pass through the v -zero cell $|s(h)|$; let this v -zero cell be denoted v_j .

Then the formula for $a(v, w)$ in $\Delta_{L(T)}$ may be written as

$$a(v, w) = \left(\prod_{i=1}^{j-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)} \right) \left(\prod_{i=j}^{n-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)} \right)$$

As in the previous cases, we have obvious equalities

$$\begin{aligned} D(v_i, v_{i+1}) &= D^h(v_i, v_{i+1}), & i \geq j \\ D(v_{i+1}, v_i) &= D_h(v_{i+1}, v_i), & i < j \end{aligned}$$

because of an isomorphism of the corresponding subspaces.

On the other hand, for $i < j$ the subspace $T(v_i, v_{i+1})$ is equal to the union of the image under ρ of $|T^h|$ and $T_h(v_i, v_{i+1})$; for $i \geq j$ the subspace $T(v_{i+1}, v_i)$ is equal to the union of the image under ρ of $|T_h|$ and $T^h(v_{i+1}, v_i)$. In both cases, prunability of T_h implies that the corresponding deviations are equal to one another. Thus

$$\prod_{i=1}^{j-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)} = \prod_{i=1}^{j-1} \frac{D_h(v_i, v_{i+1})}{1 - D_h(v_{i+1}, v_i)} = a(v, s(h)) \text{ in } \Delta_{L(T_h)}$$

and

$$\prod_{i=j}^{n-1} \frac{D(v_i, v_{i+1})}{1 - D(v_{i+1}, v_i)} = \prod_{i=j}^{n-1} \frac{D^h(v_i, v_{i+1})}{1 - D^h(v_{i+1}, v_i)} = a(|s(h)|, w) \text{ in } \Delta_{L(T^h)}.$$

This guarantees that a point v in T_h has the same barycentric coordinates when $\Delta_{L(T_h)}$ is mapped to $\Delta_{L(T)}$ as when the barycentric coordinates are computed in $\Delta_{L(T)}$ directly. \square




FIGURE 24. The inclusion of $|T_h|$ into $|T|$, the straightening of T_h in $\Delta_{L(T_h)}$, and the inclusion of $\Delta_{L(T_h)}$ into $\Delta_{L(T)}$

REFERENCES

- [Afs11] B. Afsari, *Riemannian L^p center of mass: Existence, uniqueness, and convexity*, Proc. Amer. Math. Soc. **139** (2011), 655–673.
- [Böd06] C. F. Bödigheimer, *Configuration models for moduli spaces of Riemann surfaces with boundary*, Abh. Math. Sem. Univ. Hamburg **76** (2006), 191–233.
- [CG04] R. L. Cohen and V. Godin, *A polarized view of string topology*, Topology, geometry and quantum field theory, London Math. Soc. Lecture Note Ser., vol. 308, Cambridge Univ. Press, Cambridge, 2004, pp. 127–154. MR 2079373 (2005m:55014)
- [Cha05] D. Chataur, *A bordism approach to string topology*, Int. Math. Res. Not. IMRN **2005** (2005), no. 46, 2829–2875.
- [CJ02] R. L. Cohen and J. D. S. Jones, *A homotopy theoretic realization of string topology*, Math. Ann. **324** (2002), no. 4, 773–798. MR 1942249 (2004c:55019)
- [Cos06] K. Costello, *A dual point of view on the ribbon graph decomposition of moduli space*, arXiv preprint, 2006, arXiv:math/0601130v1.
- [CS99] M. Chas and D. Sullivan, *String topology*, arXiv preprint, 1999, math.GT/9911159v1.
- [CS04] ———, *Closed string operators in topology leading to Lie bialgebras and higher string algebra*, The legacy of Niels Henrik Abel, Springer, 2004.
- [Dol95] A. Dold, *Lectures on algebraic topology*, Classics in Mathematics, vol. 200, Springer Berlin Heidelberg, 1995.
- [EK] D. Egas and A. Kupers, *Comparing combinatorial models for moduli space and their compactifications*, in preparation.
- [GK73] K. Grove and H. Karcher, *How to conjugate C^1 -close group actions*, Math. Z. **132** (1973), 11–20.
- [God04] V. Godin, *Categorical graph models in the study of the moduli space of bordered Riemann surfaces and the moduli space of smooth curves*, PhD Thesis, 2004, Stanford University.
- [God07a] ———, *Higher string topology operations*, arXiv preprint, 2007, math.AT/0711.4859v2.
- [God07b] ———, *The unstable integral homology of the mapping class groups of a surface with boundary*, Math. Ann. **337** (2007), no. 1, 15–60.
- [Har88] J. L. Harer, *The cohomology of the moduli space of curves*, Theory of Moduli (Montecatini Terme, 1985), Lecture Notes in Math., vol. 1337, Springer, Berlin, 1988, pp. 138–221.

- [Hat02] A. Hatcher, *Algebraic topology*, Cambridge U., 2002.
- [HW14] N. Hingston and N. Wahl, *Compactified string topology*, Various public lectures, 2014.
- [Igu02] K. Igusa, *Higher Franz–Reidemeister torsion*, AMS/IP Stud. Adv. Math., vol. 31, Amer. Math. Soc./International Press, Providence/Boston, 2002.
- [Iri14] K. Irie, *Transversality problems in string topology and de Rham chains*, arXiv preprint, 2014, <http://arxiv.org/abs/1404.0153>.
- [Iri15] ———, *A chain level Batalin–Vilkovisky structure in string topology and decorated cacti*, arXiv preprint, 2015, <http://arxiv.org/abs/1503.00403>.
- [Kar14] H. Karcher, *Riemannian center of mass and so called Karcher mean*, arxiv:1407.2087, 2014.
- [Kon92] M. Kontsevich, *Intersection theory on the moduli space of curves and the matrix Airy function*, Comm. Math. Phys. (1992), no. 147, 1–23.
- [Kup13] A. Kupers, *Constructing higher string operations for manifolds using radial slit configurations*, <http://math.stanford.edu/~kupers/radialslitoperationsnew.pdf>, 2013.
- [Mal11] E. Malm, *String topology and the based loop space*, arXiv preprint, 2011, <http://arxiv.org/abs/1103.6198>.
- [Pen87] R. C. Penner, *The decorated Teichmüller space of punctured surfaces*, Comm. Math. Phys. **113** (1987), no. 2, 299–339.
- [Poi10] K. Poirier, *String topology and compactified moduli spaces*, PhD Thesis, 2010, City University of New York.
- [PR11] K. Poirier and N. Rounds, *Compactifying string topology*, arXiv preprint, 2011, arXiv:1111.3635v1.
- [San12] O. Sander, *Geodesic finite elements on simplicial grids*, Int. J. Numer. Meth. Engng. **92** (2012), no. 12, 999–1025.
- [Str84] K. Strebel, *Quadratic differentials*, vol. 5, Springer-Verlag, 1984.
- [Sul05] D. Sullivan, *String topology background and present state*, Current developments in mathematics (2005), 41–88. MR 2459297 (2010c:55007)
- [Tam09] Hirotaka Tamanoi, *Stable string operations are trivial*, Int. Math. Res. Not. IMRN (2009), no. 24, 4642–4685. MR 2564371 (2010k:55015)
- [Tam10] H. Tamanoi, *Loop coproducts in string topology and triviality of higher genus TQFT operations*, J. Pure Appl. Algebra **214** (2010), no. 5, 605–615.
- [TZ06] T. Tradler and M. Zeinalian, *Algebraic string operations*, arXiv preprint, 2006, math.QA/0605770v1.

CENTER FOR GEOMETRY AND PHYSICS, INSTITUTE FOR BASIC SCIENCE (IBS), 77 CHEONGAM-RO, NAM-GU, POHANG-SI, GYEONGSANGBUK-DO, REPUBLIC OF KOREA 790-784

E-mail address: `gabriel@ibs.re.kr`

URL: `cgp.ibs.re.kr/~gabriel`

DEPARTMENT OF MATHEMATICS, NEW YORK CITY COLLEGE OF TECHNOLOGY, 300 JAY STREET, BROOKLYN, NY

E-mail address: `kpoirier@citytech.cuny.edu`

URL: `websupport1.citytech.cuny.edu/faculty/kpoirier`

REASONING MIND, 2000 BERING DRIVE, SUITE 300, HOUSTON, TX

E-mail address: `nathaniel.rounds@reasoningmind.org`